# Ecologically Compatible Autonomous AI Production

Wanhong HUANG
huangwanhong@serendip.ngo

**Abstract**

Increasingly autonomous AI-mediated production can expand task-level capabilities while leaving the renewal of machinery, energy provision, knowledge, labour, and material supply dependent on wider productive relations. This exploratory paper examines the conditions under which those capacities can be reproduced without undermining affected ecological processes. Historical-materialist analysis and a provisional generative-relational (GR) representation are used to distinguish operational autonomy, continued productive capacity, changes in critical dependencies, ecological viability, and institutional authority over feasible alternatives. Domain-specific reviews assess documented capabilities and limits in robotics, automated laboratories, production control, computing, power and cooling, supply chains, repair, and materials recovery. A typed relational model makes physical flows, enabling conditions, evidentiary provenance, and decision rights jointly inspectable without treating them as interchangeable quantities. Its incremental value relative to life-cycle assessment, material-flow analysis, reliability engineering, conventional production networks, and established provenance methods remains to be tested. Conditional formal constructions are connected to proposed counterfactual comparisons, empirical measurement protocols, and explicit identification and falsification criteria. The paper distinguishes technical feasibility, ecological effects, distributional consequences, and normative justification, allowing incompatibility or indeterminate results. It reports no original production experiment, causal ecological estimate, or demonstrated self-reproducing AI factory and establishes no general ranking of ownership forms. Its contribution is a bounded research architecture and staged empirical programme for investigating the historically constituted conditions of ecologically compatible autonomous production.

**Keywords:** historical materialism; autonomous production; productive reproduction; ecological generativity; industrial robotics; infrastructural dependence; material circulation; generative-relational inquiry.

*Discussion-draft note.* Version 3.25 is the internally consolidated release candidate of the complete 30-section exploratory manuscript. The previous adversarial review tested the scope of selected formal and empirical arguments through constructed counterexamples and explicit assumption failures. Source-level checks and formal comparison proposals do not constitute new experiments or proof of original predictive capabilities. Citation-command tracking is an indexing milestone rather than full-text validation of every cited claim: many checks are restricted to source scope, multiple citations require separate assessment, and some historical identifications remain heuristic. The manuscript presents no original factory experiment, identified causal ecological effect, universal ownership ranking, or independent peer review. Conditional formal results remain dependent on their assumptions.

*Licence and preparation.* Original material © 2026 Wanhong Huang, CC BY-NC 4.0 (https://creativecommons.org/licenses/by-nc/4.0/); third-party rights remain unaffected. ChatGPT assisted with source discovery, drafting, editing, and LaTeX preparation. The human author retains responsibility for the accuracy of claims, interpretations, and any subsequent submission.

DRAFT

## 1 Introduction

The expansion of artificial intelligence into the planning, monitoring, and execution of material production invites a renewed inquiry into the historical conditions through which societies reproduce their productive capacities. An industrial installation may carry out a growing range of operations through robotic manipulation, model-based scheduling, automated inspection, and adaptive experimental design. Yet its continued operation depends upon inputs and relations that lie beyond the immediate machine

or facility: accumulated scientific knowledge, the fabrication and replacement of components, electrical infrastructure, extraction and processing of materials, human maintenance, institutional permissions, and ecological processes supporting material supply and waste assimilation. A growing capacity to conduct productive operations without immediate human direction can therefore coexist with substantial dependence upon other workers, infrastructures, territories, and living systems. The central question of this paper concerns the reorganization of those dependencies and the conditions under which technologically autonomous production can remain compatible with ecological continuation.

This question concerns a historically specific configuration of productive forces and relations of production. In Marx's analysis of machinery, the use of scientific knowledge and mechanical power acquires its social character through the organization of production, the ownership of instruments, and the relations governing labour [120]. The *Grundrisse* further considers how accumulated social knowledge is objectified in fixed capital; the frequently discussed passage on the general intellect provides a historical resource for examining knowledge embodied in machinery, although it supplies neither a description of contemporary machine-learning models nor a determinate prognosis for them [125]. The development of automated production presents the enduring analytical problem in another technical form. Productive knowledge may be created across scientific communities, education systems, public infrastructures, industrial workplaces, and distributed software development, while the practical capacity to deploy its outputs depends upon particular arrangements of capital, equipment ownership, access rights, skilled labour, and material inputs. The relations connecting collective knowledge to effective production thus warrant investigation alongside the capabilities of the resulting machines.

The object of inquiry is *productive reproduction*: the ongoing maintenance, replacement, renewal, and transformation of the conditions that make productive activities possible. This notion is used at several distinguishable levels. First, a productive system must remain operational at the level of its assigned tasks, technical components, reliability requirements, and throughput. Second, the infrastructures that sustain those components must themselves be reproduced through energy supply, tools, trained personnel, knowledge transfer, fabrication capacity, and material circulation. Third, productive activities are organized through social relations that condition investment, labour, appropriation, distribution, and access to the means of production. Fourth, the extraction, transformation, and circulation of materials interact with ecosystems whose regenerative and evolutionary dynamics cannot simply be substituted by maintaining technological equipment. These levels are related but do not share a single metric. A system may maintain specified output while imposing pressures on ecological functions, or may exhibit improvements in material efficiency while transferring environmental burdens into distant supply chains. The analysis consequently distinguishes the reproduction of technical capacity, the reproduction of social relations, and the continuation of ecologically significant processes.

Contemporary developments give this older theoretical problem new empirical content. The International Federation of Robotics reported approximately 603,000 industrial robot installations worldwide during 2025 and an operational global stock of approximately 5.079 million industrial robots, according to its September 2026 *World Robotics* reporting [69, 70]. These figures indicate the extension of industrial robotic equipment, although an installed robot should not be counted as an independently reproducing factory. In a different domain, the International Energy Agency estimated worldwide data-centre electricity consumption at approximately 485 TWh in 2025 and projected roughly 950 TWh for 2030 in its April 2026 central outlook [68]. This projection covers data centres generally, not solely computing for autonomous material production; it should be interpreted as evidence about an important supporting infrastructure and its prospective energy demands rather than as a direct estimate of energy required for autonomous manufacturing. The IEA's 2026 critical-minerals outlook also documents continuing supply-chain concentration and highlights the relevance of strategic minor minerals to high-technology and digital systems [67]. Together these sources motivate investigation of the physical and geographical conditions of technological autonomy. They do not themselves establish that artificial intelligence caused an increase in total material extraction, nor that the observed installations possess high levels of cognitive, maintenance, or reproductive independence.

The wider social metabolism supplies an additional reason to ask how technical advances affect

DRAFT

Table 1: Selected empirical contexts for the investigation of AI-mediated productive reproduction. The indicators describe different systems; none directly measures system-wide reproductive autonomy.

| Indicator and period | Reported magnitude | Appropriate evidential interpretation |
|---|---|---|
| Industrial robots, world, 2025 | Approximately 603,000 installations; 5.079 million operating units [70] | Annual installation flow and year-end operating stock; neither measures whole-factory autonomy. |
| Data-centre electricity, world, 2025 and 2030 | 485 TWh estimated for 2025; 950 TWh projected for 2030 [68] | Supporting computational infrastructure; 2030 is a modeled projection. |
| Material extraction, world, 1970 and 2024 | Approximately 30 and 106.6 billion tonnes per year [184] | Long-run total economy-wide material throughput; no AI-specific causal attribution. |
| Electronic waste, world, 2022 | 62 million tonnes generated; 22.3% formally documented collection and recycling [101] | End-of-life conditions for all reported electronic-equipment categories. |
| Occupational exposure, world, ILO 2025 model | About 25% with some generative-AI exposure [74] | Task-based exposure, not observed unemployment or industrial automation. |

the reproduction of their material conditions. The United Nations Environment Programme's *Global Resources Outlook 2024* reports an economy-wide series rising from approximately 30 billion tonnes in 1970 to an estimated 106.6 billion tonnes for 2024 [184]. That long-period series covers economic activity as a whole and cannot be assigned to contemporary artificial intelligence. It does, however, locate autonomous production within an already extensive transformation of material throughput. The International Telecommunication Union and the United Nations Institute for Training and Research estimated that the world generated about 62 million tonnes of electronic waste in 2022, of which 22.3% by mass was documented as formally collected and recycled [101]. The measurement concerns electrical and electronic equipment broadly, not a specific class of AI systems. It nonetheless illustrates the gap between technological use and documented materials recovery. Reproduction of an automated facility involves the lifecycle of its equipment as well as its immediate operation, and claims of circularity must take account of recovery yields, reprocessing inputs, losses, secondary-material quality, and continuing primary extraction.

The relation between technical automation and living labour requires equally careful measurement. The International Labour Organization's 2025 global index estimated that roughly one quarter of worldwide employment was in occupations with some exposure to generative AI, while approximately 3.3% was in its highest exposure category [74]. This is a task-based occupational exposure estimate rather than an observed displacement rate, and it primarily concerns generative-AI capabilities rather than factory robotics alone. Such distinctions matter for any historical claim about an emergent structure of production. The technical potential to automate a task does not specify whether employers adopt the technology, reorganize work, alter working hours, intensify surveillance, retain workers for complementary functions, or displace employment through changing demand. A facility's apparent independence from direct human execution may likewise depend upon external maintenance crews, supply-chain labour, data work, repair specialists, and services supporting workers' reproduction. A historical-materialist account must follow these relations beyond the visible point of operation.

Table 1 places selected measures beside the propositions they can legitimately motivate. The table is intentionally heterogeneous. Its entries refer to different years, denominators, spatial boundaries, measurement methods, and kinds of uncertainty, so they are not components of a composite autonomy or sustainability index. The substantive empirical review in Section 5 will examine these and related series in greater methodological detail.

The relevant historical question is accordingly more specific than whether machines replace workers or whether optimization reduces emissions. It is how changes in the capacities of machinery alter the structure of necessary labour, investment, scientific knowledge, technical maintenance, supply-

DRAFT

chain dependence, and resource use, while different relations of production condition the choices made among technologically available alternatives. Under investor-owned, publicly operated, cooperative, or commons-oriented institutions, physically similar technologies can face different imperatives, resource constraints, forms of oversight, and distributions of productive benefits. None of these organizational categories guarantees an ecological outcome. Nor is technical change a mechanically sufficient cause of a particular future institutional order. The aim is to identify which material and social mechanisms mediate changes in feasible production, and how those mechanisms can be observed and tested.

The ecological dimension cannot be derived from conventional measures of plant-level efficiency alone. An autonomous production process can achieve lower energy consumption per unit of output and still increase total energy demand if the scale or composition of production changes. Equipment may be durable locally but dependent upon material extraction elsewhere; recycling may reduce demand for particular inputs while creating additional energy use, emissions, or exposure along the recovery chain. The restoration of machinery and the recovery of damaged ecosystems also exhibit distinct causal mechanisms and time scales. Work in ecological historical materialism, including analyses of the metabolic relations between social production and nature, motivates inquiry into the ways productive expansion may compromise its own biophysical conditions [41]. Yet ecological compatibility must be established in specific contexts using biophysical inventories, ecological states, exposure pathways, and recoverability evidence. The use of the term *metabolism* does not itself demonstrate a given mechanism, and a general critique of extractive production cannot replace measurements of actual impacts.

The paper's guiding distinction is between *operational autonomy* and *reproductive autonomy*. Operational autonomy concerns the extent to which a specified system can complete a defined set of productive tasks under stated service requirements without particular external interventions. Reproductive autonomy concerns its capacity to maintain or restore that operational ability over a specified horizon, accounting for maintenance, component replacement, knowledge and personnel requirements, energy and material supplies, and dependencies crossing the selected system boundary. The measures are necessarily boundary- and horizon-relative. An industrial robot may be capable of highly automated operations during its working cycle while relying on specialized human interventions or non-substitutable parts when a critical component fails. Consequently, an increase in task-level independence does not logically establish a corresponding increase in infrastructural self-reproduction. The proposed *reproduction-gap hypothesis* concerns whether this divergence is systematically important in identifiable production settings, not whether a perfectly self-sufficient industrial system is theoretically possible.

A second analytical focus is the *relocation of dependencies*. Automation may reduce certain requirements for human execution while changing exposure to chips, energy systems, communications, process equipment, upstream materials, or contractual and legal control over technology. These changes involve heterogeneous flows and permissions that resist meaningful aggregation into an undifferentiated index. The related proposal of *scarcity migration* asks whether effective production constraints shift toward alternative inputs or institutional interfaces when previously costly operations become less demanding. Historical and contemporary research on input–output relationships, industrial ecology, supply-chain resilience, maintenance systems, and economic power already address significant portions of this terrain. The intended contribution of a generative-relational (GR) representation is narrower: to test whether explicitly recording typed dependencies, their evolution, their controlling institutions, and the opportunities they enable or foreclose adds identifiable explanatory or predictive value to competent established approaches. Its relevance depends upon comparative performance; a representational framework does not acquire independent empirical authority through its conceptual breadth.

The investigation is organized around four connected questions. *First*, what activities, support interventions, equipment-restoration obligations, boundaries, and horizons must be specified before claims of operational or reproductive autonomy are comparable? *Second*, which observed changes in required inputs, substitution options, maintenance burdens, and effective access rights explain the relocation of productive dependencies and associated transformations of labour? *Third*, under what measured conditions can a defined productive service be sustained while ecological states remain within independently justified and scientifically defensible constraints? *Fourth*, how do financing, information, authorization, and

DRAFT

governance alter the set of implementable production policies and the allocation of their consequences? The questions follow different evidentiary routes: activity and repair records for autonomy, supplier and access histories for dependency changes, ecological observations for continuation, and decision and institutional records for implementability. An answer to one cannot substitute for the others.

Four proposed mechanisms organize the subsequent tests: a possible gap between task execution and replenishment of productive capacity; relocation rather than simple elimination of necessary inputs; conditional compatibility between technical continuation and site-specific ecological continuation; and institutional selection among technologically feasible production pathways. These mechanisms can operate together without moving in the same direction. A facility might automate inspection, diversify component supply, reduce absolute life-cycle burdens, and nevertheless remain dependent on a non-substitutable proprietary repair service. Another may gain technical independence while increasing pressures on a vulnerable watershed. The paper therefore asks which changes can be identified separately, and which alternative explanations survive comparison with established industrial-ecological, economic, and reliability methods. A credible result may be favorable, adverse, null, or unidentified; none is to be made confirming merely by redescribing it as a relational transformation.

The study develops these questions through three connected modes of inquiry. A historical-materialist analysis examines the development of productive forces within changing relations of appropriation, labour, knowledge, and social reproduction, without treating technological change as a self-sufficient historical agent. A domain-specific technical review evaluates industrial robotics, autonomous scientific laboratories, production control, computing infrastructures, energy systems, supply chains, repair, and material recovery, assigning weight to evidence according to the scope and maturity of each demonstration. For example, the corrected account of the A-Lab inorganic-materials experiment reports a bounded 17-day campaign realizing 36 of 57 targets; this is evidence for an experimentally specific form of adaptive laboratory autonomy, with an explicit post-publication correction, not for an autonomous industrial system capable of reproducing its laboratory infrastructure [178, 179]. Finally, a comparative methodology investigates matched production policies, dependency networks, environmental inventories, labour conditions, and institutional settings, including feasible counterfactuals and conditions that would weaken the proposed theoretical synthesis.

The conceptual architecture draws provisionally on GR inquiry developed in related work on ecological knowledge formation through oceanic robotics. Here the material production of artefacts, the continuation of productive activities, and the reproduction of their enabling infrastructures constitute distinct objects from ecological observation. The transfer is thus conceptual and methodological. It does not assume that a formalism proposed for tracing scientific observations is automatically adequate for supply chains, social institutions, ecological dynamics, or industrial control. Where standard tools, including life-cycle assessment, material-flow accounting, reliability analysis, input–output methods, and recorded data provenance, are sufficient, no additional GR vocabulary is required. Where additional representation is proposed, the paper specifies prospective comparisons against those competent alternatives. Value categories likewise remain distinct: monetary exchange, productive use, scientific benefit, ecological continuation, and institutional legitimacy may be related without constituting quantities naturally reducible to one another.

The argument proceeds through thirty numbered sections, with *section* as its highest organizational level. Sections 2–5 develop the historical analysis and contemporary empirical setting. Sections 6–12 examine typed production relations, autonomy, recursive dependencies, bottlenecks, appropriation, and the reproduction of labouring capacities. Sections 13–20 review domain-specific engineering, energy, supply, repair, and materials-recovery evidence. Sections 21–25 investigate ecological continuation, recoverability, life-cycle consequences, robust viability, and cumulative scale effects. Sections 26–29 then address institutionally feasible choices, causal policy comparison, dependency measurement, and inferential limits. Section 30 synthesizes the resulting claims and specifies a staged research programme.

The present inquiry is a theoretical and research-design investigation. It reports no original factory deployment, new robot trial, primary life-cycle inventory, causal labour estimate, or independent ecological field evaluation. Its historical interpretations, proposed GR constructs, and formal feasibility models

DRAFT

are therefore presented with distinguishable evidential status. The immediate scholarly task is to turn the historical question of productive reproduction into empirically discriminating comparisons: between operational and reproductive capacity, between local efficiencies and system-wide outcomes, between technical feasibility and ecological viability, and between alternative institutional rules for selecting and sustaining productive activity. The next section establishes the historical-materialist categories needed for these comparisons and addresses their scope and competing interpretations.

# 2 The Historical Organization of Productive Reproduction

The historical organization of productive reproduction establishes the explanatory foundations for the paper's account of autonomous AI production. The central question concerns how societies maintain and transform the capacity to produce, rather than the productivity of a machine considered in isolation. It requires attention to the material conditions under which production recurs, the social relations that organize access and appropriation, the activities that reproduce human capacities, and the ecological processes on which production depends. Historical materialism offers an account of the connection among these processes, although competing interpretations differ over their causal priority and mechanisms. This section reconstructs the relevant categories, distinguishes their explanatory levels, and sets out the conditions under which they may guide empirical inquiry without presupposing a determinate future social formation.

## 2.1 Material Production Within Historically Constituted Relations

In the 1859 preface to *A Contribution to the Critique of Political Economy*, Marx presents the relations entered into through social production alongside the development of material productive forces, arguing that contradictions among them are central to historical transformation [119]. These concepts define an investigation into the organization of social existence. Productive forces include labour capacities, instruments, methods of coordination, and usable knowledge; they operate through particular arrangements of access, control, obligation, and cooperation. Relations of production concern the socially constituted positions through which persons and institutions direct, sustain, and appropriate productive activity. Property and legal arrangements are consequential dimensions of those relations, although their practical operation also depends on enforcement, bargaining power, and infrastructure.

This formulation does not render technical artifacts inert. A power loom, a numerically controlled machine, or an automated production cell can change feasible tasks and the organization of work. Yet a machine's incorporation into social production depends upon investment, labour organization, material supply, skills, and the purposes governing its operation. The same physical type of equipment may be operated within an investor-owned company, a public research establishment, a cooperative, or a shared production facility. Those differences affect whose needs guide output, the treatment of maintenance and ecological costs, the ability to change productive decisions, and the distribution of resulting capacities. Conversely, institutions face constraints arising from actual production possibilities. The relation is reciprocal, historically situated, and mediated rather than reducible to a general opposition between technology and society.

A useful distinction concerns the *capacity* to perform an operation and the *social possibility* of exercising that capacity. An industrial robot might be capable of assembling a component given specified materials and programming. Whether the component can be manufactured repeatedly also depends on continued parts supply, tool wear, available energy, authorized access, compatible software, skilled intervention, and financing. Technically feasible operations thus form only one subset of the trajectories available to a historically organized production system. This section uses this distinction provisionally; the production-boundary and dependency representations developed in later sections will make its system assumptions explicit.

Historically constituted relations should also be distinguished from an assumption that a society is governed by one homogeneous mode of activity. Household subsistence, public provision, contractual

DRAFT

employment, cooperatives, informal labour, and commercial exchange can coexist and become interdependent. Their significance may differ across regions, industries, and institutional scales. A historical-materialist inquiry must accordingly specify which arrangements are being examined, whose actions they organize, and how their material reproduction is achieved. The presence of AI-mediated coordination alone does not justify classifying a production system as a new mode of production.

## 2.2 The Reproduction of Production Beyond the Individual Establishment

Using chapter numbering from the linked English editions (which differs across some historical editions), Marx develops a decisive distinction in the analysis of *simple reproduction* in *Capital*, Volume I: a productive activity viewed as a recurrent social process must replenish the conditions of its continuation, while repetition under capitalist relations also contributes to reproducing those relations [121]. His treatment of the reproduction of aggregate social capital in Volume II examines the relationship between production of means of production and production of consumption goods, including the replacement of used-up productive inputs [124]. These analyses prevent an individual enterprise's sales, revenue, or short-run efficiency from being confused with society-wide continuity.

An isolated productive operation can consume inventories, degrade machinery, or draw down ecological stocks while reporting a successful output. Its repetition requires replacement inputs and continuing services, many of which originate beyond the establishment. In Marx's two-department reproduction analysis, the material forms of outputs matter alongside their value magnitudes: a monetary amount sufficient to purchase replacement equipment cannot itself supply equipment that has not been produced and made accessible. The scheme is not a detailed input–output account of modern industry, but it makes analytically visible the intersectoral conditions of reproduction [124]. At contemporary industrial scale, these conditions include electrical systems, software maintenance, calibration equipment, communications networks, upstream fabrication, and industrial waste management.

Reproduction can preserve an existing production scale or participate in expanded accumulation. Marx examines the reinvestment of surplus value into additional productive capacity under capitalist conditions and emphasizes that expansion requires material means of production and means of subsistence rather than a merely financial decision [122]. This distinction remains relevant when automated systems support higher output or lower unit costs. Profitable reinvestment, industrial capacity extension, and adequate replacement of physically necessary inputs are different propositions. The expansion of one facility may also alter the conditions of reproduction for suppliers, displaced workers, public infrastructure, and ecological systems. At a societal level, a growing aggregate of productive installations could coexist with the deterioration of some of its enabling conditions.

We therefore adopt *productive reproduction* as a relational research object: the activities and dependencies through which productive capability, necessary inputs, and institutional access are sustained or transformed across a specified boundary and assessment horizon. A system's future capacity cannot be inferred solely from current throughput, even when its performance is highly automated. An observation period can show that productive activity continues with limited intervention; a longer or differently bounded assessment may reveal reliance on external repairs, replacement tools, trained personnel, or cumulative depletion. This definition is proposed for the present inquiry and does not claim to supersede the diverse historical and economic meanings of reproduction in Marxian scholarship.

## 2.3 Accumulation and the Reproduction of Social Relations

In a capitalist production process, reproduction concerns both material means and the historically specific relation through which labour power, productive assets, and surplus are mobilized. Marx's distinction between simple reproduction and accumulation clarifies how recurrent production may renew the social conditions of wage labour and capital ownership, while the conversion of part of surplus value into additional capital can change the scale of production [121, 122]. This account cannot be transposed without qualification to a publicly funded laboratory, cooperative workshop, or commons-maintained computational resource. Such arrangements likewise require material replacement and social coordination, but

DRAFT

may employ different rules of appropriation, contribution, investment, and distribution.

For the present paper, the analytical sequence is to identify a technology's actual productive capability, reconstruct the social arrangements necessary for its operation, and then examine the reproduction of those arrangements through production. These processes should be studied separately before their interaction is assessed. A privately controlled automation platform may be maintained through public scientific education, publicly funded research, user-generated data, supplier expertise, or contracted maintenance. The institutional recognition and monetary remuneration of these contributions can differ from their material importance. Likewise, a public production facility may depend on proprietary operating software or externally controlled fabrication equipment. Public ownership of a facility does not by itself establish independent access to the infrastructure required for continued operation.

*Appropriation* here means the historically organized processes through which actors obtain or control claims over outputs, productive means, knowledge, income, or decision authority. Those objects differ in legal form and measurability. The ownership of a machine, the permission to use a software model, the power to purchase scarce inputs, the right to reproduce a process, and the authority to decide which ecological impacts are acceptable may be assigned to different actors. They should therefore be traced independently. The distinction matters for evaluating AI production because automated coordination can modify some operational tasks while leaving control over financing, equipment, intellectual property, and market access unchanged or altered in different directions.

Capital accumulation is one historically specific mechanism for reproducing and extending productive capacity. The broader investigation must also encompass arrangements in which surpluses are allocated through budgets, membership decisions, commons rules, or planned provisioning. A comparison cannot infer ecological or distributive consequences from a regime's name: governance effectiveness, investment conditions, public accountability, compliance, and technical feasibility must be observed. Conversely, technical performance cannot establish the legitimacy of a regime of appropriation. The paper consequently examines the interaction between technical possibilities and institutional reproduction without presenting either as sufficient to determine the other.

### 2.4 The Reproduction of Labouring and Social Capacities

Productive continuity requires the reproduction of human capacities, including sustenance, learning, health, care, collective coordination, and the transfer of knowledge. In Marx's treatment of social reproduction, the maintenance of labour power is connected to historically constituted consumption requirements and the reproduction of capitalist relations [121]. Subsequent social-reproduction scholarship broadens the investigation to activities and institutions that sustain everyday life, including care, education, households, public services, and unevenly distributed forms of unpaid work. Fraser argues that a comprehensive account of capitalism needs to examine social reproduction, political power, and ecological conditions as enabling and contested conditions of economic production [43]. These dimensions should not be treated as a peripheral supplement to a machine-centred model.

AI-mediated automation makes the distinction between *task substitution* and *social reproduction* especially important. A system may replace a specific assembly, classification, or scheduling operation while continuing to depend on human development of technical methods, construction and servicing of equipment, supervision, workplace coordination, medical and educational systems, and social institutions through which workers live. Changes in the labour required at an observed facility are not equivalent to changes in the total labour needed across its supply network. Neither do reduced labour hours directly establish improvements in employment security, wages, autonomy, or the distribution of necessary care work. These are historical and empirical relationships requiring separate evidence.

A further distinction concerns the valuation and recognition of such activities. Labour contributing to technological capacity may be remunerated through wages, contracts, public funding, or indirect social provision, while other essential activities remain unpaid or minimally recorded in organizational accounts. Whether a transition to autonomous production reduces, intensifies, relocates, or reorganizes labour depends partly on market demand, the availability of complementary skills, contractual relations,

DRAFT

Table 2: Analytical domains of productive reproduction and their distinct evidential requirements. The categories overlap causally but do not form interchangeable measures of success.

| Domain | Question of continuation | Evidence needed for a bounded assessment |
|---|---|---|
| Technical capacity | Can productive functions be maintained or restored? | Equipment lifetime, repairability, replacement inputs, availability, service dependencies and operating horizon. |
| Institutional relations | Can recognized rights, obligations and authority remain operative or change? | Ownership and licensing records, procurement constraints, governance decisions, access conditions and effective enforcement. |
| Social reproduction | How are human capacities and necessary supporting activities sustained? | Working conditions, labour time across supply chains, care arrangements, skill formation, income and public provision. |
| Ecological continuation | Which ecological processes remain viable following productive intervention? | Spatially explicit material and energy flows, exposure, ecological response, recovery evidence, uncertainty and relevant baselines. |

and wider institutions. Section 12 will treat those mechanisms in detail. At this stage it is sufficient to establish that the reproduction of labouring and social capacities cannot be inferred from the apparent independence of an automated machine.

## 2.5 Ecological Conditions of Productive Continuation

Historical production also operates through material exchanges with ecological systems. Natural processes supply biophysical conditions for agriculture, extraction, energy generation, water use, atmospheric regulation, and material transformation. Industrial infrastructures can alter these conditions through land conversion, pollution, resource depletion, and cumulative effects. Such interactions enter a historical-materialist account of reproduction both as constraints on productive activity and as transformations of ecological relations whose significance extends beyond human economic use.

Ecological interpretations of Marx differ in emphasis, but Foster's reconstruction of the *metabolic rift* identifies an important intellectual precedent: capitalist organization can disrupt material exchanges and regenerative processes on which social production depends [40]. Saito's study of Marx's ecological writings likewise provides an account of the significance of natural-scientific inquiry to Marx's critique of political economy [162]. These contributions support taking material ecological interactions seriously while leaving open the specific causal effects of contemporary AI systems, which require separate measurement. A production arrangement may aggravate, mitigate, or reorganize ecological burdens depending on technology, scale, regulation, and patterns of demand.

For this paper it is necessary to distinguish at least three forms of continuation: the reproduction of *technical productive capacity*, the reproduction of *social and institutional arrangements*, and the continuation or regeneration of *ecological processes*. They influence one another but need not move in parallel. Replacement components may restore the operation of a production line while extraction to manufacture those components contributes to habitat disturbance elsewhere. A financially self-sustaining enterprise may shift material burdens onto communities or ecosystems beyond its accounting boundary. Conversely, restoration and remanufacturing practices could improve some ecological conditions while increasing near-term costs or maintenance requirements. Each case requires specified boundaries, measured effects, and an explicit counterfactual.

Table 2 specifies a methodological separation needed for the later comparative analysis. A single financial, energetic, or productivity indicator cannot stand in for all four domains. Even within an ecological assessment, resource efficiency and ecological viability have different evidential objects. A lower energy requirement per unit of output may be offset by increased aggregate throughput or by disturbance channels not represented in energy statistics. The social importance of ecological continuation also cannot be derived entirely from its instrumental contribution to industrial output. The selection of acceptable

DRAFT

risks, habitat losses, or compensatory actions involves ecological knowledge and legitimate normative procedures, examined in Sections 21–29.

## 2.6 Historical Explanation and the Limits of Technological Determination

An important methodological dispute concerns the priority of productive forces relative to relations of production. Cohen's analytic reconstruction offers a systematic defence of an explanation assigning primacy to the development of productive forces [24]. In contrast, Brenner's work on agrarian class structures emphasizes historically specific class relations in accounting for divergent paths of economic development [15]. Wood likewise argues against accounts in which capitalism emerges as an inevitable extension of trade, stressing the historically particular social relations and imperatives of capitalist accumulation [192]. These are significant differences in explanatory emphasis, and this paper does not settle them by adopting AI technological change as a privileged historical cause.

A further warning against a simplified one-way causal account comes from Engels's later clarification that the material production and reproduction of life should not be equated with a claim that economic variables are the sole determinants of historical outcomes; political and legal forms may materially influence the shape of historical development [32]. The present research treats technology, labour, institutional authority, and biophysical constraints as interacting causal candidates whose effects must be established in a given historical setting. This approach is compatible with taking material limits seriously while avoiding the presumption that one technology necessarily generates one relation of ownership or one ecological trajectory.

Methodologically, the dispute calls for distinguishing three questions. The *origin* of an institution or technical configuration concerns how it arose under particular historical conditions. Its *reproduction* concerns which mechanisms sustain it across observed production episodes. Its *transformation* concerns the pathways through which those mechanisms change or fail. A capability introduced by machine learning might be adopted because of a labour shortage, a procurement programme, a competitive strategy, or a scientific objective. Once installed, its continuation may depend on patents, energy contracts, interoperable equipment, or public infrastructure that were less salient during its development. Its effects on labour, output, and ecology then depend on subsequent decisions and conditions. The three questions have overlapping evidence but are not interchangeable.

This distinction also constrains counterfactual inference. Correlations between the adoption of automation and the concentration of ownership cannot determine whether the technology caused concentration, whether concentrated firms were more capable of adopting the technology, or whether an unobserved institutional condition shaped both. Claims about ecological compatibility likewise require comparisons against materially plausible production alternatives, including situations where a less autonomous configuration performs better. Historical explanation consequently benefits from comparative cases, reconstructed decision sequences, matched technical baselines, and explicit uncertainty about unobserved causal mechanisms.

## 2.7 Implications for Autonomous AI Production

The preceding reconstruction gives autonomous production a more precise historical research object. Operational automation concerns capabilities realized within selected productive episodes. Productive reproduction concerns the maintenance, replacement, and transformation of the relations and conditions that make those episodes possible. Historical materialism demands that both be studied through the material organization of production, the relation between collective capacities and control, and the distribution of the burdens of reproduction. The ecological extension adds an examination of continuing biophysical processes that may be altered or diminished by the very interventions sustaining technical capacity.

A concrete research comparison can therefore begin with two facilities capable of producing the same class of output at equivalent quality. If one relies on more autonomous execution, the relevant inquiry extends to equipment longevity, external replacement and maintenance, upstream material and

DRAFT

energy use, labouring conditions across suppliers, rights over necessary software, and measured ecological effects. The outcome could be that automation reduces some dependencies, diversifies repair options, and lowers total burdens. It could instead concentrate reliance on a non-substitutable component or relocate pollution to an upstream production site. Both outcomes remain admissible under the historical framework developed here. Ownership or governance comparisons require the same care: different institutional arrangements may enlarge or restrict feasible production in ways that are not predictable from their labels.

This section establishes four premises for the subsequent investigation. First, productive forces are realized within historically specific relations of access and appropriation. Second, the continuity of production requires replenishing material, social, and institutional conditions beyond the immediate operation. Third, technical, social, and ecological reproduction are analytically distinct, even where their causal interactions are strong. Fourth, historical transformations cannot be derived from technical capability alone; they require inquiry into mechanisms, institutions, ecological constraints, and viable alternatives. Section 3 will examine a specific historical development within this architecture: the incorporation of collectively generated knowledge into machinery and the changing terms under which that knowledge becomes effective productive capacity.

# 3 The Incorporation of Collective Knowledge into Machinery

The historical reproduction of production examined in Section 2 depends upon more than the replacement of equipment and material inputs. Productive capabilities also incorporate accumulated knowledge: practical techniques, scientific theories, technical standards, records of experimental failure, methods of coordination, and the expertise needed to maintain and revise them. This section investigates how such knowledge becomes materially effective through machinery, computation, and the organization of labour. Its purpose is to distinguish the collective origins of productive intelligence from the historically variable arrangements through which particular actors gain access to, govern, and appropriate its application. Marx's analysis of fixed capital provides the principal historical-materialist point of departure; subsequent histories of industrial automation and contemporary evidence on research, AI development, and intellectual property clarify both the continued relevance and the limits of this analogy.

## 3.1 Social Knowledge as a Constituent of Productive Capacity

A machine does not embody knowledge merely because an expert designed it. Its operation presupposes a historically accumulated body of possibilities for acting on materials, transmitting forces, measuring change, coordinating sequences, and addressing failure. A cutting tool depends on metallurgical knowledge and manufacturing tolerances; an industrial controller depends on mathematical representations, sensors, electronics, software protocols, and records of operation; a machine-learning-guided laboratory additionally depends on training data, an experimental ontology, laboratory procedures, reliable instruments, and criteria for recognizing a successful result. These resources are developed and reproduced through multiple institutions and activities rather than through a single, self-contained act of invention.

Three distinctions are consequential. First, *knowledge production* involves the formation, evaluation, correction, and circulation of representations and practical competences. Second, *productive incorporation* occurs when selected competences and representations become causally effective through a particular instrument or coordinated labour process. Third, *productive command* concerns the rules and effective powers through which the instrument can be used, modified, maintained, or withheld. These are related processes, but their institutional subjects, temporalities, and material requirements can differ. The publication of a method does not itself fabricate a working device. Conversely, possession of an operating device does not establish the capacity to recover its design knowledge or to replace components once they wear out.

The historical relation is reciprocal. Practical production generates observations, anomalies, skills, and demands that condition subsequent scientific work. Laboratories and engineering organizations con-

DRAFT

vert some of this experience into transferable representations; those representations, when embedded in tools, can reorganize production and create further occasions for learning. Knowledge may thus be productive before it is encoded in a machine, and its future development may depend on access to the material systems in which it is exercised. A historical-materialist interpretation must account for this circulation rather than describe innovation as a unilateral transfer of intellectual discoveries into otherwise passive industry.

### 3.2 The Objectification of Knowledge in Marx's Analysis of Machinery

The passages on machines and fixed capital in Marx's *Grundrisse*, drafted in 1857–1858 but published posthumously and cited here through a later English translation, describe the growing significance of general social knowledge within the material powers of production [125]. Machinery appears as a historically produced objectification of collective capacities, within which science and cooperation acquire productive effectiveness. In this argument, the concept frequently described as *general intellect* concerns the societal development of knowledge and productive organization. It is neither an anticipation of a particular machine-learning architecture nor a claim that intellectual work resides entirely inside a physical apparatus. The objectification of knowledge is a social and material process: instruments, techniques, organization, and labour become mutually implicated in production.

In *Capital*, Volume I, Marx locates machinery within the organization of the factory and examines its effects on labour, the labour process, and the capitalist pursuit of surplus value [120]. His account is analytically different from celebrating technological progress in the abstract. The incorporation of science into an organized system of machinery can increase productivity while altering workers' practical discretion, the distribution of tasks, supervision, and control over the labour process. The intellectual capacities embodied in the system may confront particular workers as an external power where ownership and direction of productive assets are separated from the workers who operate them. The relevant contrast concerns historically specific social relations, rather than an inherent moral property of scientific knowledge or automation.

The two texts should also be distinguished. The *Grundrisse* notebooks constitute an exploratory manuscript; *Capital* develops a more elaborated critique of the organization of production under capitalist relations. The term *general intellect* should therefore be used with interpretive restraint. Extending it to AI requires a reconstructed argument about how social knowledge is produced, materially incorporated, and institutionally organized, rather than an assertion that contemporary AI is itself Marx's general intellect. Nor does the category supply an empirical measure of how much collective knowledge resides in a given model or machine.

For the present research, the useful historical proposition is narrower and testable at the level of mechanism: productive intelligence can be socially generated, materially embodied, and institutionally controlled through different arrangements. The interests and capacities that guide the technical design of machinery may differ from those of the persons whose activity generated its antecedent knowledge or who depend on its operation. The subsequent questions are which contributions matter to the machine's operation, how they can be recognized, and how access to the resultant capacity is organized.

### 3.3 The Social Selection of Technical Designs

Studies of industrial automation have demonstrated that the architecture of productive machinery is shaped by decisions concerning work organization and authority. David Noble's history of numerical control describes competing approaches to automating machine tools and argues that institutional investment, military procurement, and managerial preferences influenced which approach acquired momentum [135]. Its significance for this paper lies in the possibility that a chosen technology may reflect a particular allocation of control among workers, engineers, managers, and outside institutions. Technical feasibility alone does not explain why one arrangement of knowledge and productive authority is selected over another.

DRAFT

This historical finding should not be elevated into a universal explanation. Technical reliability, production volume, workplace skill, market competition, and available infrastructure can also influence a design decision. Some techniques that separate planning from execution may deliver valuable consistency and safety; others may impair operator judgement or create brittle dependence on unavailable experts. Comparative investigation must reconstruct the actual alternatives, the constraints under which they were assessed, and the evidence for their performance. It should allow cases in which decentralization improves adaptability and cases in which centralized coordination reduces risk or waste.

Recent scholarship on the social history of AI further emphasizes the relevance of labour organization, classification, and historically accumulated patterns of perception and coordination to computational intelligence [150]. Such interpretations are valuable interlocutors rather than exhaustive accounts of AI's technical development. The empirical architectures of contemporary learning systems require analysis of statistical estimation, optimization, computing hardware, experimental practices, and particular deployment settings. An explanation grounded exclusively in labour organization would leave important mathematical and physical mechanisms unresolved, just as an explanation grounded exclusively in algorithms would understate their institutional conditions.

We can accordingly investigate a *design-selection mechanism*: available techniques define a set of candidate production architectures; material costs, organizational authority, compatibility, skilled labour, regulatory constraints, and investment choices help determine which architecture is developed and deployed. The historical conclusion concerns contingent selection among feasible alternatives. The further GR question is whether those alternatives can be represented as changing sets of productive possibilities under distinct access and control relations. That formal question belongs to Sections 6–10; the present section establishes the historical basis for posing it.

### 3.4 The Material Translation of Research into Productive Equipment

Between published research and an operating production system lie multiple technical and institutional interfaces. A method may be conceptually clear yet lack robust training data, production-grade sensors, a qualified controller, validated tooling, or a maintenance pathway. The conversion of knowledge into effective productive capacity involves practical work at these interfaces, with uncertainty and failure at each stage. An exploratory representation of this sequence is: formation and critical testing of knowledge; development of reproducible procedures; engineering embodiment; integration with an operational environment; and maintenance, correction, and reproduction of the resulting capability. This is a conceptual sequence of functions, not a universal chronology in which one phase has to end before another begins.

These interfaces are especially important in AI-mediated production. A learning algorithm can be publicly described, while the model that performs adequately in a particular facility depends on site-specific calibration, sensors, compatible actuators, operational data, material-quality specifications, human judgement, and institutional permission. A model's reported task performance can therefore coexist with a large, incompletely documented infrastructure of enabling conditions. A facility can also possess the technical means to perform an action without having the right to obtain the data or materials required for that action, or without adequate evidence of safety in the intended environment.

We distinguish four evidentially different outcomes of incorporation. A representation may be *available* for inspection; a technical implementation may be *reproducible* within a declared environment; an installation may be *operationally effective* on a specified task; and the resulting capacity may be *reproductively durable* under maintenance demands, disturbances, and material dependencies. Evidence at one interface should not be presumed to establish success at the next. For example, open availability of model weights does not imply access to fabrication equipment, energy, industrial data, or authorization to use a production facility. This is one way in which the social circulation of knowledge and the physical reproduction of production can diverge.

A further distinction concerns the attribution of contributions. Scientific understanding may emerge from public laboratories and universities; implementations may be developed by commercial engineers;

DRAFT

Table 3: Selected indicators of the contemporary organization of knowledge and AI development. These measures represent different populations and cannot be interpreted as one accounting identity or as direct evidence of appropriation.

| Domain and year | Reported observation | Interpretive limit |
|---|---|---|
| US academic R&D, FY 2024 | US$118 billion total expenditure; 55% federally funded; 63% classified as basic research [131] | US higher-education sector; categories describe expenditure and research type, not ownership of subsequent inventions or AI-specific activity. |
| Notable AI models worldwide, 2025 | 91.58% attributed to industry by the 2026 AI Index's sector classification [174] | Curated subset of "notable" models; organizational attribution cannot quantify the origins of training data, labour, scientific methods, or deployed industrial value. |
| GenAI patent families worldwide, 2025 | 37,808 published families, compared with 18,862 in 2024 [194] | Published patent families are an innovation/IP activity indicator, not granted patent counts, infringement findings, productive deployment, or social welfare measures. |
| Data work in high-stakes AI, 2021 | Interview study of 53 practitioners reports recurrent data-quality cascades [163] | Qualitative sample across selected professional settings; it does not estimate the global incidence of undervalued labour or consequences in manufacturing. |

process data may be generated by workers or equipment; deployment may depend on suppliers and municipal utilities. These contributions cannot be assigned to one actor merely because an application is marketed or controlled by that actor. Nor does the existence of diverse antecedent contributions independently establish an unlawful appropriation or a particular distributional entitlement. Such conclusions require legal, contractual, and normative analysis of the actual case. The historical object is a plurality of participating activities whose relevance may be obscured by a narrow accounting of the final artifact.

## 3.5 Contemporary Evidence on Research and Control over AI Development

Current evidence allows this historical problem to be specified more concretely, although no single indicator measures the embodiment or appropriation of collective intelligence. Table 3 juxtaposes measurements from different segments of the knowledge-to-production trajectory. The distinct dates, populations, and measurement conventions are integral to their interpretation.

The US National Center for Science and Engineering Statistics reports that higher-education institutions performed approximately US$118 billion of research and development in fiscal year 2024, with the federal government accounting for 55% of that expenditure and basic research representing 63% of the activity [131]. This documents a substantial publicly funded knowledge-producing infrastructure in one country. It does not trace any particular industrial model to those expenditures, quantify how much of a model's value originated in universities, or demonstrate that a given enterprise appropriated uncompensated public research. Such claims would require study of actual knowledge lineages, contracts, licensing, procurement, and deployment.

On another measurement boundary, the 2026 Stanford AI Index, drawing on Epoch AI's compilation, assigns approximately 91.58% of globally recorded *notable AI models* in 2025 to industry organizations [174]. The result suggests that the resources required for developing models within this selected frontier set are increasingly organized in industry. The denominator is neither all AI applications nor all machine learning research. Publicly supported research, academic contributions, collaborations, and open methods may enter those models' developmental histories, but their magnitudes and causal significance cannot be inferred from the sectoral attribution of the final models.

A further institutional interface is observable in patent activity. The World Intellectual Property Organization's July 2026 technology update identifies 37,808 published generative-AI patent families for 2025, compared with 18,862 in 2024 [194]. Patent publication indicates claimed invention activity within a selected technological classification. It is not a count of patents granted or enforceable, and it is not equivalent to an operational factory, a socially productive innovation, or a demonstrated increase in ecological burdens. The combination of public research finance, industry-associated frontier models, and increasing patent-family publications motivates an investigation into how knowledge crosses institu-

DRAFT

tional boundaries; the aggregate statistics themselves do not identify the relevant transfers or determine their legitimacy.

These limitations are methodologically productive. A credible historical-materialist account should reconstruct particular pathways connecting publicly accessible research, proprietary development, industrial deployment, and practical control. It should also seek counterexamples: cases of open models generating broad productive access, shared public infrastructure supporting robust local industry, licensed collaboration with fair attribution, and industrial advances whose crucial knowledge emerged within private research facilities. The objective is to explain variation in the incorporation of knowledge, not to interpret every commercialized machine through a prespecified account of extraction.

### 3.6 The Labour Required to Materialize and Sustain AI Systems

Treating AI as embodied knowledge does not permit the disappearance of the labour required to generate, evaluate, and maintain it. Collective intellectual and practical contributions include designing instruments, collecting and cleaning data, interpreting anomalies, annotating examples, repairing sensors, monitoring failures, managing versions, and translating model outputs into actions. These activities differ in technical content, skill, employment relation, geographical location, and degree of visibility to managers or users. Their relation to AI-mediated production cannot be inferred solely from the number of people who operate the final model interface.

An interview study by Sambasivan and colleagues, involving 53 AI practitioners working in selected settings across India, countries in East and West Africa, and the United States, documents failures propagated by insufficient attention to data practices in high-stakes AI projects [163]. Its qualitative findings are pertinent to the organizational status of data work, even though the cases are not representative of all industrial AI or all national labour markets. The study supports a mechanism of practical dependence on often under-recognized upstream activities. It does not establish that all data labour is underpaid, that all AI systems exhibit the same failures, or that a single remedy is suitable across institutions.

The distinction between *task performance* and *productive reproduction* becomes especially clear here. A model may perform a specified prediction while requiring continuous observation, data correction, physical calibration, and human decisions about exceptional conditions. If those activities are omitted from an account of autonomy, apparent improvements may reflect a reclassification or displacement of labour rather than its removal. Conversely, improved instrumentation, appropriate automation, and better data governance may reduce burdensome work and increase workers' capacity to exercise informed discretion. An empirical account should measure task changes, responsibility, skill development, occupational safety, and maintenance burdens alongside output and error rates.

The distribution of productive intelligence also raises epistemic questions. A worker or technician may possess local, partly tacit knowledge of a process without control over the technical documentation or model updates that formalize it. When this knowledge is encoded into production systems, some aspects become transferable and reproducible; other dimensions may be simplified, lost, or rendered difficult to revise. A comparable issue arises when field observations enter datasets while the people able to identify changing conditions cannot contest the dataset's categories. These possibilities require case-specific evidence, not an assumption that explicit encoding always impoverishes practical knowledge. They help explain why the provenance and revision history of productive representations could matter to both reliability and the distribution of authority.

### 3.7 Institutional Appropriation and Access to Productive Intelligence

The incorporation of knowledge into machinery becomes economically consequential through historically specific rules of access and control. Relevant instruments include ownership of equipment, trade secrecy, patent rights, copyright where applicable, contractual restrictions, technical standards, model licensing, access to compute and data, and the practical ability to maintain an installation. They are analytically heterogeneous. Legal ownership of a machine may permit its ordinary use while access to the software, replacement components, diagnostic records, or model parameters required for modification

DRAFT

remains limited. An open scientific method may coexist with non-transferable assets such as specialized instruments or scarce production materials.

The historical-materialist question is how these different forms of command govern the realization of productive possibilities. Ownership and intellectual-property rights can help finance research, establish duties, support licensing, and structure collaboration; they can also produce exclusion or dependence under specified conditions. Public and commons-oriented arrangements can widen access, yet they must still finance maintenance, manage security, ensure accountability, and respond to scarcity. Cooperative or state ownership likewise leaves open questions of actual participation, downstream dependency, and ecological oversight. The effects of an institutional form are therefore empirical and historically contingent rather than deducible from its label.

An adequate theory of heterogeneous value circulation should also recognize that economic returns, scientific value, technical reliability, workers' capacities, and ecological continuation are different evaluative objects. The generative perspective on value creation developed in Ron Eglash's *Generative Justice* offers a relevant intellectual inspiration for considering how productive benefits are generated and circulated [31]. The formal GR framework proposed in this paper remains an independent and provisional investigation; it neither follows from that source nor presumes that these diverse values share a commensurable numerical measure.

Two counterfactual configurations illustrate the problem. In one hypothetical facility, workers retain access to process data and diagnostic controls while a publicly financed design is maintained through a cooperative service network. In a second, an externally maintained platform provides reliable operation but prohibits local modification, leaving repair dependent on a single vendor. Either arrangement could yield greater throughput in particular conditions. Their difference would become evident when a component fails, an ecological constraint changes, an operator discovers an unsafe condition, or a replacement procedure becomes necessary. The empirical study should compare the costs and benefits of each arrangement, the effective rights to correct failures, the accessibility of substitute infrastructure, and the resulting feasible responses. Neither ownership model is assumed to determine the preferred outcome in advance.

### 3.8 Knowledge Infrastructures and Their Ecological Conditions

The material embodiment of productive intelligence depends on ecological and physical infrastructures. Scientific laboratories require instruments and controlled environments; data centres require energy, electronic components, cooling, and replacement hardware; industrial robots depend on fabricated materials, electricity, logistics, and the disposal or recovery of failed equipment. Some knowledge can be copied at low marginal computational cost, yet the systems through which it is operationalized have material lifetimes, finite energy demands, and geographically uneven resource dependencies. A digitally transferable model therefore does not imply a freely reproducible material production capacity.

This observation has consequences for the historical-materialist treatment of collective knowledge. The capacity to make a prediction, devise a chemical route, or schedule a machine can grow without a proportional improvement in material replacement capacity or ecological viability. A factory may gain knowledge about resource-efficient operation and nevertheless increase its aggregate resource demand through expanded output. Alternatively, better scientific understanding can support reduced losses, longer component lifetimes, and improved ecological monitoring. Identifying which trajectory occurs requires a specified system boundary, baseline production function, ecological indicators, and a counterfactual comparison. Ecological compatibility cannot be inferred directly from the sophistication or accessibility of knowledge.

The same issue applies to the reproduction of knowledge infrastructures themselves. Long-term scientific capability requires retaining and revising records, cultivating skills, supplying instruments, preserving public access where appropriate, and allowing critical scrutiny of methods. A proprietary model with high observed performance may have a fragile knowledge-reproduction pathway if independent verification is impossible or specialist maintenance is unavailable. An open model may have a different

DRAFT

vulnerability if durable funding, reproducible hardware, or ecological cost accounting is absent. The historical formation of knowledge as a productive force thus links directly to the problem of reproducing productive capacity addressed in Section 2.

### 3.9 Research Implications for AI-Mediated Productive Reproduction

The preceding discussion supplies a set of mechanisms for further investigation. Social knowledge is generated through heterogeneous labour and institutions; selected representations are translated into operating machinery; design alternatives are chosen within particular technical and social constraints; rights and infrastructures affect access to productive capabilities; and the conditions supporting those capabilities require ongoing reproduction. These propositions do not establish an invariant historical law. They specify relations whose strength and direction can differ across production settings.

Four empirically criticizable hypotheses follow. *First*, facilities incorporating externally developed models or methods may exhibit a difference between the organizations credited with operational performance and those contributing essential upstream scientific, technical, and data resources; this difference can be investigated through provenance, finance, contracts, and maintenance histories. *Second*, technology selections may change the distribution of authority and technical discretion independently of measurable productivity; comparative design histories and workplace records can identify whether such changes occurred. *Third*, greater availability of information or algorithms may leave reproductive autonomy unchanged where access to physical infrastructure or repair knowledge remains constrained; matched implementation cases can test this divergence. *Fourth*, broadened access to productive knowledge can support ecological improvement, increased throughput and burdens, or both; lifecycle and production-policy comparisons are required to distinguish them.

Each hypothesis admits competing explanations and meaningful negative findings. Measured productivity differences might be attributable to site-specific equipment rather than the organizational locus of knowledge. Changes in authority might be negligible where labour institutions preserve practical discretion. The availability of standardized components and open maintenance documentation might allow operational and reproductive autonomy to improve together. Expanded information sharing might reduce aggregate ecological burdens even after changes in production scale are included. These outcomes would refine the proposed explanation rather than constitute a failure of the entire historical-materialist inquiry.

The immediate methodological requirement is to connect historical interpretations to observable entities and events. Relevant records include research provenance, technical standards, model and software dependencies, design alternatives, ownership and licensing instruments, workforce tasks, maintenance histories, material flows, and ecological measurements. Section 4 will examine the historical development of industrial automation through this lens, while later formal sections will specify how such heterogeneous dependencies might be represented without confusing documented associations, physical causal effects, and normative entitlement. The contribution of the present section is to establish the incorporation of collective knowledge into machinery as an empirically situated transformation of productive possibilities, whose results depend on both technological embodiment and the relations organizing its continuation.

## 4 The Historical Development of Industrial Automation

The incorporation of collective knowledge into machinery, examined in Section 3, has taken different material forms as the organization of production has changed. This section reconstructs selected transitions in industrial automation to identify how the capacity to execute, program, coordinate, and adapt productive activity has been redistributed among workers, machines, firms, and infrastructures. The historical question concerns the conditions that made particular architectures of automation feasible and attractive, the alternatives that coexisted with them, and the relations of production they enabled or altered. The discussion compares mechanized factory organization, electrification and moving assembly, numerical

DRAFT

Table 4: Selected historical changes in the organization of industrial automation. Periods denote examples and diffusion histories rather than exclusive technological epochs.

| Illustrative development | Expanded productive capability | New or reorganized conditions of dependence |
|---|---|---|
| Mechanized factory systems (eighteenth–nineteenth centuries) | Repeated operations driven through machines and collective power transmission | Capital equipment, motive power, maintenance, factory supervision, labour discipline |
| Electrified mass production (late nineteenth–twentieth centuries) | Reconfigurable mechanical drive and paced flow across specialized stations | Electrical infrastructure, standardized components, throughput coordination, upstream inputs |
| Numerical control (1950s onward) | Tool motion specified through encoded geometric instructions | Programming, metrology, reliable controls, machine-tool service, procurement |
| Industrial robots and programmable controllers (1960s onward) | Repeatable manipulation and changeable industrial sequences | Integration, tooling, sensing, safety, system programming, repair |
| Networked and model-supported production (late twentieth century onward) | Coordination of inventory, design, quality, and production across facilities | Software interoperability, communication, data quality, global suppliers |
| AI-assisted and closed-loop systems (twenty-first century) | Bounded adaptation of planning or experimentation using new observations | Model validation, computing, data, independent assessment, material replenishment |

control, industrial robotics, programmable coordination, and recent AI-supported production. These developments overlap in time and differ substantially among industries and regions; they are treated as historically contingent configurations rather than successive stages through which every productive system must pass.

A distinction guides the comparison. *Task execution* refers to the physical performance of an operation; *process specification* to the encoding of an operation and its conditions; *coordination* to the synchronization of heterogeneous activities; and *revision* to changing these activities in response to observations, failures, or new objectives. Technical progress in one dimension does not establish autonomy in the others. A machine may execute a program repeatably while depending on externally produced instructions, maintainers, supply chains, and institutional decisions. Table 4 identifies selected innovations and their distinct evidential implications.

## 4.1 Mechanization and the Reorganization of the Labour Process

Early industrial mechanization reorganized how mechanical force and human activity were combined. In textile mills, machine tools, mining, and other sectors, production increasingly relied on fixed installations and sequences of specialized operations. The resulting factory arrangements changed the spatial and temporal relations among workers, materials, power sources, and supervisory functions. The power of machinery was inseparable from the physical infrastructure supplying motion, the capacity to maintain equipment, and access to raw materials and markets. A production process could therefore exhibit considerable mechanization while retaining extensive human work in feeding, supervising, adjusting, inspecting, and repairing machines.

Marx's analysis of the factory system identifies the organization of machinery as one means by which socially developed productive powers are incorporated into a coordinated labour process [120]. His account illuminates the possibility that the increased collective effectiveness of labour can coincide with changes in control over its execution. At the same time, the historical consequences of mechanization should not be inferred from the machine alone. Patterns of industrial finance, skills, labour bargaining, energy supply, access to markets, and legal relations differed across sectors. A critical historical expla-

DRAFT

nation asks which changes followed directly from a machine's technical properties and which arose from the historically specific institutions governing investment and work.

Mechanization accordingly presents an initial version of the reproductive problem. The factory could concentrate labour and increase output while generating new requirements for coal or water power, replacement components, infrastructure maintenance, transport, and regular access to materials. The boundary of apparently self-contained machine production expanded outward to encompass the networks that supplied and reproduced those conditions. The same relation will recur, in technically different form, in computationally coordinated production.

## 4.2 Factory Electrification and the Organization of Mass Production

Electrification enlarged the ways in which mechanical power could be supplied and distributed within a factory. Its significance extended beyond substituting electric motors for steam or other prime movers. The capacity to power equipment more flexibly allowed alterations in plant layout, machinery arrangement, process sequencing, and maintenance. Paul David's historical comparison of the dynamo and computer stresses the long interval between the technical availability of electrical power and its wider productivity consequences, partly because existing factories embodied earlier investments and organizational designs [27]. A recent establishment-level study of North Carolina manufacturers during 1905–1926 likewise finds productivity and output gains associated with electrification, alongside mixed distributional consequences: higher average wages, larger returns to skill, and a lower labour share in the studied setting [26]. Those findings apply to the examined establishments and historical identification strategy; they are neither a global estimate nor a general law governing the distributional effects of electricity.

The moving automobile assembly line illustrates how machinery, component standardization, supply timing, and the detailed organization of work can develop together. The Ford Motor Company's Highland Park moving line was introduced in 1913, drawing on earlier techniques from multiple industries rather than appearing as a single invention without antecedents [113, 136]. Its industrial significance lay in coordinated flow: a sequence of activities was linked so that the movement of partly completed products governed the timing of work. The arrangement could reduce assembly time while raising the importance of line balance, interchangeability, inventory reliability, and supervision. The substantial productivity gains associated with this reorganization do not, by themselves, determine how the gains were allocated or how workers experienced the resulting forms of coordination.

David Nye's historical study emphasizes electrification and the combination of existing technical elements in the assembly line, complicating accounts that attribute the new organization solely to scientific management [136]. The significance for the present paper is methodological. A machine's performance and the organization of production form a coupled historical process. Investment in machinery may require changes in layout, labour training, standardization, and institutional coordination before its technical possibilities become materially effective. Later digital technologies exhibit analogous complementarities, although the technologies, periods, and forms of evidence differ.

## 4.3 The Programmability of Machine Tools

Numerical control introduced a distinctive separation between the physical motion of a machine tool and the encoded specification of its path. The development associated with John Parsons, the Massachusetts Institute of Technology's Servomechanisms Laboratory, and US Air Force support culminated in a working numerically controlled milling-machine demonstration in 1952 [128, 127]. Instructions encoded on punched media could direct operations that previously depended more heavily on particular combinations of manual positioning, templates, and machine-specific arrangements. This capability made certain geometrically complex operations reproducible from transferable descriptions, but it also required the production of reliable programs, functioning controls, calibrated equipment, and effective translation between geometric design and machine motion.

The history of numerical control is especially important because technical feasibility and institutional selection were interwoven. David Noble reconstructs the development and diffusion of numerical con-

DRAFT

trol as a history of engineering choices shaped partly by military procurement, firm organization, and the distribution of authority between programmers, managers, and skilled machinists [135]. In his interpretation, alternatives maintaining greater shop-floor discretion were sometimes available, and control relations helped influence the direction of investment. This is a substantive historical argument requiring engagement with competing explanations, including precision requirements, production scale, changing product geometry, machine reliability, and the costs of preparing control instructions. Its value here is to establish a research question about design selection, not to presume that every programmable machine was introduced to diminish worker autonomy.

Programmability therefore created two analytically separate capabilities. It made the execution of some operations less dependent on an operator's continuous manual positioning; and it allowed the instructions for these operations to be stored, communicated, modified, and reproduced outside the machine's immediate physical location. Yet programming still depended on knowledge, testing, interpretive judgement, and machine-specific expertise. The numerical-control example shows why the transfer of an instruction into software cannot be equated with the autonomous reproduction of the machinery, knowledge, and industrial organization required to execute it.

## 4.4 Industrial Robotics and Reprogrammable Manipulation

Industrial robotics extended programmed execution to handling, transfer, welding, painting, and other manipulative operations. The Unimate installed in a General Motors facility in Trenton, New Jersey, in 1961, for unloading a die-casting press is an instructive early example documented in a preserved industrial artifact record [181, 71]. The early robot relied on an externally specified sequence and substantial mechanical and control infrastructure. Its achievement was repeatable performance of a hazardous industrial operation, rather than independent perception, planning, maintenance, or resource acquisition. Subsequent developments in electromechanical actuation, sensor feedback, computing, and tooling widened the range of tasks and operating conditions [71].

These architectures also reveal the difference between flexibility and generality. A robot arm may be reprogrammable across several operations while requiring different end effectors, fixtures, object presentations, safety systems, and validation procedures for each task. Manufacturing integration depends upon the behaviour of the whole production cell. The reliability and economics of an installed system therefore cannot be derived from the nominal degrees of freedom of its manipulator or from isolated laboratory demonstrations. Later sections examine these engineering constraints directly; the historical role of robotics here is to show that programmable physical work increases organizational options while generating new requirements for integration and maintenance.

Robotic diffusion has also been highly uneven. According to the International Federation of Robotics' September 2026 release covering calendar year 2025, factories worldwide installed approximately 603,000 industrial robots that year, while the estimated stock in operation exceeded five million units [69]. These are different measures: annual additions describe a flow, while operating stock concerns accumulated installed equipment. The reported stock does not measure the number of autonomous factories, the degree of decision autonomy of installed machines, or the distribution of benefits among workers and owners. Its significance for this historical section is to demonstrate the contemporary scale of industrial robotic incorporation without attributing all such adoption to AI.

## 4.5 Programmable Coordination Across the Production System

The historical expansion of automation also involved the control of production sequences outside individual machine tools and manipulators. Programmable logic controllers emerged in the late 1960s partly in response to the cost and inflexibility of revising hard-wired relay systems. The development of the Modicon controller and General Motors' interest in reconfigurable industrial control show how machine-level programmability became part of broader industrial coordination [165, 10]. These devices did not constitute autonomous industrial agents; their significance was the capacity to change the logic linking sensors, machines, and sequences without rebuilding the same physical switching arrangements

DRAFT

for every product or process revision. The cited company history and retrospective engineering account establish the origin and architectural purpose of programmable control; neither is a comparative dataset on global adoption or the distribution of its benefits.

Subsequent developments in computer-aided design, manufacturing software, inventory management, and networked logistics integrated more activities into data-mediated industrial organization. Such systems can support coordination across separate facilities, but they also make compatibility, communication infrastructure, trusted records, specialized software, and upstream suppliers more consequential. Production dependencies may become less visible to observers restricted to an individual factory. A low level of manual intervention at the point of assembly may depend on considerable human, technical, and material activity elsewhere in the system.

The diffusion of networked production also requires an institutional explanation. Firms select between internal provision and external procurement, integrated planning and modular subcontracting, centralized standards and locally maintained systems. Public infrastructure, international trade, industry standards, and labour institutions condition these choices. At the same time, practical maintenance and situated work can correct or bypass formal computer representations. A historical reconstruction should therefore include failures, workarounds, and organizational revisions as well as successful deployment narratives. Apparent gains in coordination can involve new vulnerability to outages, incompatible interfaces, concentration of suppliers, and reduced observability of distant material burdens.

### 4.6 Contemporary AI and the Boundaries of Adaptive Production

Machine learning and AI-supported control can add new kinds of inference to established programmable and networked production. Under specified conditions, systems may use sensor histories to revise quality predictions, adapt process parameters, schedule maintenance, select candidate experiments, or compare simulated alternatives. Unlike a fixed programmed sequence, an adaptive policy may depend on observations whose interpretation is statistically uncertain. Its reliability consequently depends on training and evaluation populations, physical instrumentation, validation against the deployed environment, and safeguards against model error. The coexistence of learning systems with conventional controllers is therefore an engineering architecture to investigate, rather than evidence of their wholesale replacement.

Autonomous scientific laboratories provide a bounded illustration. The A-Lab study reported a 17-day experimental campaign involving 353 experiments and realization of 36 of 57 targeted inorganic compounds through an integrated combination of computation, machine learning, robotics, experimental measurement, and active learning [178, 179]. The cited report, subsequently associated with an author correction, describes a particular materials-synthesis workflow. Its results should be interpreted within that experimental setting; they do not demonstrate the autonomous production of replacement laboratory instruments, consumables, energy, or all necessary research objectives. Nor does a model's success in proposing experiments imply that its inferred phase composition or projected ecological consequences will be correct in a different setting.

Digital models offer another interface between production and prediction. Reviews of production digital twins document both potential uses in factory planning and continuing difficulties concerning definitions, interoperability, data quality, implementation, and verification [109, 12]. A digital representation becomes scientifically useful when its scope and model discrepancies are known, observations can be connected to well-defined physical quantities, and its decisions can be checked against competent baselines. Verification of a digital representation against its implementation and validation of its predictions against observations are distinct empirical tasks. Simulation success is one grade of evidence; reliable field performance and sustained operational reproduction are others. The addition of AI to digital twins does not erase these distinctions.

The 2026 Stanford AI Index identifies substantial progress in selected AI benchmarks alongside limitations in physical-world robotics; such assessments depend on particular task sets and evaluation protocols [174]. This juxtaposition counsels against interpreting increasing general-purpose model performance as direct evidence of generalized industrial autonomy. The material capabilities of AI-mediated

DRAFT

production should be specified at the level of particular operations, sensors, actuators, constraints, and time horizons.

### 4.7 Productivity, Labour, and Competing Historical Explanations

Different historical explanations attach different causal significance to machinery. A technological-progress account may emphasize growing capabilities and diffusion of more effective techniques. A historical-materialist account additionally asks how capital accumulation, class relations, work organization, and institutional control shape the adoption and use of those techniques. Labour-process scholarship, including the interpretation of deskilling associated with Harry Braverman, addresses changes in managerial control and the organization of labour; subsequent research has emphasized that skill effects and worker discretion vary among technologies and contexts [170, 135]. These approaches provide hypotheses for comparing organizational outcomes, not permission to infer the motives of a particular firm without evidence.

Empirical research on robots and labour demonstrates why the evidence must be interpreted at the appropriate scale. Graetz and Michaels, using industry-level data from seventeen countries for 1993–2007, estimate a contribution of approximately 0.36 percentage points to annual labour-productivity growth from increased robot use, report no statistically significant reduction in aggregate employment within their industry-country panel, and identify a decline in the low-skilled employment share [50]. Acemoglu and Restrepo, using variation in robot exposure across US commuting zones, estimate negative effects on employment and wages in more exposed local labour markets [3]. These findings concern distinct outcomes, spatial units, identification strategies, and observation periods. One does not cancel the other; they jointly demonstrate that aggregate production gains can coexist with unequal local or occupational consequences.

The historical question is therefore how particular production arrangements mediate the allocation of gains and losses. Employment totals alone do not capture changes in work intensity, bargaining power, skill formation, maintenance obligations, hazardous exposure, or the distribution of income and control. Conversely, a documented adverse effect in one setting cannot establish the necessity of identical effects under different institutions and technologies. This paper will investigate different feasible technical configurations and forms of ownership without presuming which arrangement is inherently productive or ecologically beneficial.

### 4.8 Material Dependencies and the Historical Interpretation of Autonomy

Across these historical changes, the productive system has acquired new capacities to perform operations and coordinate activities beyond continuous human execution. Yet the reproduction of those capacities requires accumulated machinery, spare parts, energy, training, maintenance, social knowledge, and environmental inputs. Factory electrification depended on new energy distribution and complementary investment; numerical control on programming and precise machine tools; industrial robotics on integration and servicing; and AI-supported production on computation, data, physical validation, and supply continuity. In each case a capability gained at one interface could increase the importance of another dependency.

This does not support a simple claim that automation always deepens dependence. A modular design may reduce reliance on one supplier; adaptive maintenance may lower waste or extend machine life; improved control may reduce particular environmental loads. The empirical issue concerns the change in *which* dependencies matter, their substitutability, and their distribution across the productive network. Ecological effects also require investigation beyond the factory boundary. Improved energy efficiency at the point of operation can coexist with increased throughput or material extraction, whereas some technical changes may reduce both direct and indirect burdens. Neither historical materialism nor the proposed GR framework determines those outcomes in advance.

A comparative interpretation of this history suggests three research propositions. First, increases in execution or planning capability require separate evidence from claims of infrastructure-reproductive

DRAFT

autonomy. Second, the technical relocation of labour and coordination tasks can alter the position of critical material and institutional dependencies, with consequences conditioned by the relevant production relations. Third, any claim that a historically developed production arrangement is ecologically compatible requires measurement of resource use, disturbance, recovery, and displaced burdens across declared spatial and temporal boundaries. These are candidate mechanisms for later analysis, rather than historical laws established by the selected examples.

The role of the present section has been to clarify the varied pathways through which mechanization, programming, robotics, and AI have reorganized production. The next section examines contemporary material evidence more systematically, comparing the latest indicators for industrial robots, computation, energy, materials, and labour while distinguishing observed historical changes from projections and causal hypotheses. The formal representation of changing productive dependencies then follows in Section 6.

# 5 The Contemporary Material Basis of AI-Mediated Production

The preceding historical analysis established that changes in machinery, programmability, and productive coordination depend upon the material and institutional arrangements through which technical capabilities are implemented. This section examines the contemporary empirical conditions in which AI-mediated production is developing. Its purpose is to identify the measurable scale, uneven distribution, and reproductive requirements of today's productive infrastructures, and to establish a defensible baseline for the relational and ecological models developed in subsequent sections. The available evidence covers industrial robotics, computational electricity demand, capital expenditure, strategic material processing, economy-wide extraction, equipment disposal, and exposure of occupations to generative AI. None of these indicators alone measures autonomous industrial production. Their conjunction is analytically useful insofar as each illuminates a different condition whose reproduction can become consequential when the execution and coordination of production change.

The evidence is necessarily heterogeneous. Global industrial robot installations are an annual flow of specified machines, whereas the stock of robots in use is an accumulated quantity. Data-centre electricity estimates cover computational activities beyond AI or industrial production, whereas the energy use of a particular robotic facility requires separate measurement. Expenditure measures denote financial commitments rather than physical output. Material-extraction figures represent the world economy, while electronic-waste figures refer to particular discarded product categories. Occupational exposure indexes identify task characteristics, not observed displacement. The reporting years also differ, and projections must be separated from historical estimates. Tables 5 and 6 preserve these distinctions so that historically suggestive co-movement is not mistaken for an identified causal relationship.

## 5.1 The Installed Base of Industrial Automation

The International Federation of Robotics (IFR) reports approximately 603,000 industrial robot installations worldwide in 2025, an increase of about 11% on its preceding-year comparison, and an estimated operating stock of 5.079 million industrial robots [69, 70]. The two values describe different aspects of the production system. Installations indicate the flow of new equipment entering productive use during the reporting year; the operating stock records accumulated equipment remaining in use, reflecting earlier investment and the survival, replacement, and retirement of machines. A ratio between annual installations and stock could describe their relative magnitude at a particular date but would neither establish the average service life of robots nor directly identify replacement needs. Such inferences require consistent stock-flow accounting, industry-specific lifetimes, retirement data, and reporting conventions.

Deployment is geographically uneven. According to IFR's September 2026 release, China installed roughly 354,000 industrial robots in 2025, accounting for approximately 59% of worldwide installations; the United States installed nearly 38,500 and Japan 36,219 [69]. These are annual installation counts, not robot density measures, total manufacturing investment, or the share of global robotic equipment pro-

DRAFT

Table 5: Selected indicators of expanding productive and computational infrastructure. Numbers describe different populations and are not additive.

| Indicator and source | Value, place, and reference period | Measurement status and interpretive limit |
|---|---|---|
| Industrial robots installed [69, 70] | About 603,000 units; world, 2025 | Annual installation flow; not a count of autonomous factories |
| Operating industrial robots [70] | 5.079 million units; world, end of 2025 | Estimated operating stock; distinct from installations |
| Industrial robot installations in China [69] | 354,000 units; 59% of world installations, 2025 | Location of deployment; not ownership, technological origin, or control |
| Data-centre electricity demand [68] | 485 TWh; world, 2025 | IEA estimate of all data centres, including non-AI uses |
| Data-centre electricity outlook [68] | 950 TWh; world, 2030 | Conditional central projection published in 2026; not observed demand |
| Capital expenditure by five large technology companies [68] | Above USD 400 billion; combined, 2025 | Company-level capital spending; not exclusively AI or data centres |

duced in each country. Differences in industrial composition, manufacturing scale, wages, integration capacity, production demand, and investment conditions complicate comparisons. It would be inappropriate, for example, to conclude that a country with a larger annual installation flow has achieved greater operational autonomy per unit of output, or that it controls the entire industrial knowledge and component supply chain supporting the installed machines.

For historical materialism, these indicators are informative because they show how investment in fixed productive equipment varies within a globally differentiated division of labour. Machines embody prior research, design, fabrication, and labour; their use depends upon facilities, qualified operators, tooling, electricity, integration services, software, and component replacement. At the same time, the geography of installation differs from the geography of semiconductor production, mechanical-component fabrication, system ownership, and intellectual-property control. Studying those geographically distinct relations is necessary to understand how the apparent independence of an automated plant can coexist with continuing, and sometimes concentrated, dependence on external conditions of production. This is a hypothesis about the organization of dependencies rather than a finding demonstrated by installation data alone.

Robot counts have another limitation: an industrial robot may be installed in a line whose scheduling, material feeding, quality inspection, exception management, and maintenance remain intensively human-directed. Conversely, a facility can automate planning or process optimization through software without substantially increasing its robot count. The contemporary expansion of industrial robotics therefore documents one material component of automation while leaving the extent of cross-process AI coordination and the capacity for self-maintenance unmeasured. Later sections will distinguish observed mechanical operation, control-loop autonomy, and reproduction of the machinery and infrastructure on which the system relies.

## 5.2 The Energetic Conditions of Expanding Computation

AI-mediated production depends upon computation, but computational demand cannot be read straightforwardly from robot deployment. Some controllers operate locally with limited processing requirements; other applications involve simulation, digital twins, automated experimental design, centralized model training, cloud inference, or extensive data storage. The International Energy Agency (IEA), in its April 2026 assessment, estimates worldwide data-centre electricity consumption at approximately 485 terawatt-hours (TWh) in 2025 and projects 950 TWh in 2030 in its central scenario [68]. The latter is a projection conditioned upon expected infrastructure development, adoption, bottlenecks, and efficiency. Its ratio to the 2025 estimate describes a scenario trajectory, not a measured change already realized.

DRAFT

Data centres serve multiple computational functions. Their aggregate electricity consumption should therefore not be attributed entirely to artificial intelligence, much less to AI used in physical production. Even among AI-related workloads, the allocation between model training, inference, storage, networking, cooling, and local industrial computation varies by architecture and accounting boundary. Direct measurement of the electricity consequences of a particular AI-enabled production policy would require a functional unit, such as electricity consumption per validated component or per reliably completed task, together with the electricity used in associated sensing, production machinery, communications, and upstream computation. Comparing a conventional factory with an automated alternative while counting servers for only one side would introduce systematic boundary bias.

Energetic reproduction also involves capacity and timing. Electricity measured in TWh is an amount consumed over a period; its availability at a productive site depends upon the temporal distribution of demand, local grid capacity, generation resources, transmission constraints, storage, prices, maintenance, and interruptions. A system may operate efficiently under stable electricity supply while failing to preserve output when power quality deteriorates or when repairs and grid upgrades are delayed. The IEA's updated report discusses bottlenecks associated with electricity infrastructure, grid connections, advanced chips, manufacturing capacity, and financing [68]. These constraints help explain why the growth of computational capability can be mediated by slower-moving physical institutions and networks. They do not prove that a specific bottleneck governs every region or production activity.

A more adequate account should distinguish operational electricity, embedded energy in equipment and construction, indirect supply-chain requirements, and the ecological consequences of the generation mix and local water-management practices. The intensity and direction of these burdens depend on siting, cooling and power technologies, facility utilization, replacement cycles, and policy choices. Consequently, the same software improvement could reduce the electricity intensity of an operation, expand the number of operations economically feasible, or produce both effects at different scales. Estimating the net consequence requires explicit counterfactuals and demand assumptions. The empirical claim warranted at this stage is narrower: the reproduction of contemporary AI infrastructure has a measurable energetic basis, and its scale cannot be assessed from algorithmic performance alone.

### 5.3 Capital Mobilization and the Construction of Productive Capacity

The material expansion of AI infrastructure is also an investment process. The IEA reports that the combined capital expenditure of five large technology companies exceeded USD 400 billion in 2025, with a further substantial increase estimated for 2026 [68]. The reported magnitude concerns the companies' capital expenditure and should not be relabelled as exclusively spending on AI model training, data centres, or automated manufacturing. The categories of investment include heterogeneous equipment, buildings, networks, and other capital assets, whose exact inclusion and treatment depend on the underlying financial accounts and IEA aggregation. An economy-wide inference about the profitability or ecological efficiency of AI production cannot be drawn directly from this company-level total.

Nevertheless, the figure identifies a feature of the period that matters for historical-materialist inquiry: computational capabilities are being materially reproduced through large-scale commitments of accumulated capital. Financing arrangements, depreciation rules, expected returns, access to land and power, construction lead times, procurement agreements, and technological obsolescence can shape which installations are built and which remain economically viable. A model improvement may change the economic value of existing hardware or the expected lifespan of prospective investments without changing the physical lifetime of equipment. Conversely, replacement decisions can accelerate material throughput even if older machinery remains technically functional. These processes connect the turnover of fixed capital with technical change and the material metabolism of production.

Ownership of capital assets should also be distinguished from control over operational decisions and access to productive services. A firm may rent computational capacity, rely on contracted equipment maintenance, purchase electricity through intermediaries, or license proprietary control software. The distribution of claims over revenue, maintenance responsibilities, model access, physical property, and

DRAFT

rights to interrupt service may therefore differ across the productive network. Such distinctions matter because the expansion of collective computational and engineering capacities can be compatible with different configurations of ownership and control. Later comparison among investor-owned, public, co-operative, and commons-oriented arrangements should explain these configurations without presuming that any institutional form automatically yields ecological compatibility or social benefit.

## 5.4 The Concentration of Material Processing

Advanced productive equipment contains heterogeneous material inputs whose availability depends upon extraction, refining, specialized fabrication, logistics, and recycling. The IEA's *Global Critical Minerals Outlook 2025* reports that, across its selected key energy minerals, the average market share of the three largest refining countries rose from approximately 82% in 2020 to 86% in 2024 [65]. The statistic concerns a defined mineral basket and refining stage. It cannot be generalized to every material required for semiconductor fabrication, electric motors, computing equipment, or robotics; nor is it an average share of the three largest companies. It captures geographical concentration at a particular production interface.

Concentration can matter because productive continuity depends on the substitutability of inputs, inventories, refining processes, transport, technical qualifications, and the speed with which alternative suppliers can be developed. Yet concentration alone is not a measure of coercive power, supply unreliability, or the probability of interruption. A highly concentrated supply chain might be reliable under some conditions; an apparently diversified chain could still depend upon a shared energy source, specialized precursor, or licensing arrangement. The empirical question is whether critical inputs have feasible substitutes, whether failures propagate, and which actors can modify the conditions of access.

The distinction becomes especially important when evaluating claims of autonomous material production. Even a production system that manufactures a high proportion of its own mechanical parts can depend on geographically distant processing of specialized materials, precision metrology equipment, or electronic components. A higher measured share of in-house fabrication could accompany a larger dependence on a narrow external component class. To establish this mechanism, later analysis must identify the actual bill of materials, failure-critical parts, replacement intervals, qualified supplier sets, and the costs and time needed to restore productive capability. The IEA concentration estimate motivates that examination without substituting for it.

## 5.5 Economy-Wide Extraction and the Material Conditions of Continuation

An account confined to specialized minerals would omit the broader material reproduction of economic activity. The United Nations Environment Programme's *Global Resources Outlook 2024*, prepared with the International Resource Panel, presents a historical series in which global material extraction increases from approximately 30 billion tonnes in 1970 to 106.6 billion tonnes in its 2024 reporting horizon [184]. The latter figure is a report estimate, and the series concerns fossil fuels, biomass, metallic ores, and non-metallic minerals across the world economy. It does not measure materials consumed by AI installations alone, nor establish that automation caused the growth observed across the preceding decades. Aggregate extraction also requires careful treatment of accounting methods, embodied imports, stock accumulation, and changing population and output.

For the argument developed here, the significance of this long-run series is relational: expanding productive forces remain embedded within large material flows, including flows required to build factories, energy systems, transport infrastructure, housing, and the equipment that maintains them. A robot capable of performing one production task autonomously does not thereby reproduce the infrastructure for extracting, transporting, and processing its inputs. Furthermore, the reproduction of ecological capacities cannot be inferred from an accounting balance showing that enough usable material remains available. Extraction, land conversion, waste, and emissions can alter ecological processes through mechanisms not reducible to an economy-wide mass total. Later ecological analysis will therefore examine the type, site, and timing of interventions and the capacity of affected systems to continue and recover.

DRAFT

Table 6: Material and labour indicators used to delimit the empirical research object. Indicators differ in geography, measurement method, and definition.

| Indicator and source | Value, reference period, and geography | Evidential scope and exclusion |
|---|---|---|
| Top-three refining countries [65] | Mean share 86%; selected energy minerals, world, 2024 | Geographical share of refining in a defined basket; not every AI input |
| Material extraction [184] | About 30 Gt in 1970; 106.6 Gt for the 2024 report horizon, world | Economy-wide resource extraction; 2024 endpoint is report estimate |
| Electronic waste generated [101] | 62 Mt; world, 2022 | Electrical/electronic equipment waste; not AI hardware alone |
| Formally collected and recycled electronic waste [101] | 22.3% of generated e-waste by mass; world, 2022 | Documented formal treatment; not material-input circularity |
| Circular material use [37] | 12.2% of overall material use; EU, 2024 | Secondary-input share; regional, all covered materials |
| Employment with some GenAI exposure [74] | About one quarter of employment; world, 2025 study | Occupational-task exposure index; not observed displacement |
| Highest GenAI exposure gradient [74] | 3.3% of global employment; 2025 study | High-exposure classification; not likelihood of unemployment |

Distinguishing *technical turnover* from *ecological turnover* is particularly important. Components may be replaced in weeks or years, while ecological responses to extraction or disposal can extend over different temporal horizons. High substitutability within a manufacturing inventory need not imply substitutability of ecological functions in a habitat. These claims should be treated as a programme for site-specific testing, not as a universal ranking of natural and engineered processes. What the global materials series establishes is the need to locate autonomous production inside economy-wide physical reproduction rather than imagine it as an informational process disconnected from material throughput.

## 5.6 Material Recovery and the Limits of Circularity

The capacity to recover materials changes the reproduction of industrial systems, but its empirical meaning depends on what is measured. The International Telecommunication Union and United Nations Institute for Training and Research report that approximately 62 million tonnes of electronic waste were generated worldwide in 2022, of which 22.3% was documented as formally collected and recycled in an environmentally sound manner [101]. The denominator is globally generated electrical and electronic equipment waste under the report's classification. The numerator describes documented formal collection and recycling, not a universal estimate of every recovery activity, and the value cannot be interpreted as the proportion of all industrial materials successfully recovered. It is also two years earlier than the European circular-material indicator below.

Eurostat reports that 12.2% of material use in the European Union came from recycled inputs in 2024 [37]. This *circular material use rate* describes the contribution of recovered materials to a region's material input. It has a different numerator, denominator, geography, year, and physical scope from the global electronic-waste collection statistic. Directly comparing 12.2% with 22.3% as though they were competing measures of the same phenomenon would be methodologically misleading. Both figures, however, indicate that material recovery is a partial and institutionally mediated process, with large differences among product categories and economic regions.

Technological material recovery also involves quality, energy consumption, disassembly, hazardous constituents, material mixing, requalification, and the feasibility of reusing recovered components in their original applications. Recovering a metal by mass need not restore the particular purity, form, or component functionality required for a semiconductor, sensor, or actuator. Conversely, reuse, remanufacturing, repair, and extended equipment lifetimes can preserve productive capacities without requiring material recovery through melting or chemical processing. A research design focused only on recycling

DRAFT

percentages would overlook these distinct pathways of reproductive continuity.

The economic and ecological consequences depend on the displaced primary inputs and the additional burdens incurred by recovery. System-level evaluation must therefore measure the quality of recovered outputs, processes and energy used in recovery, transport, losses, contamination, and changes in equipment lifetime. Subsequent sections will use industrial-ecological methods to test whether material recovery reduces aggregate burdens under declared functional units and boundaries. This empirical question remains open for any specific AI-mediated production architecture.

## 5.7 Task Exposure and the Reorganization of Living Labour

Employment indicators illuminate a further dimension of productive reproduction: the socially constituted capacities of workers who develop, operate, supervise, maintain, and live alongside productive systems. The International Labour Organization's 2025 refined index finds that approximately one quarter of global employment falls in occupations with some exposure to generative AI, and 3.3% falls within the index's highest exposure gradient [74]. Its construction combines task classifications, surveys, expert assessments, and AI-supported scoring. The values identify potential changes in the content of work; they are not estimates of workers dismissed, hours actually automated, wages changed, or productive facilities that function without people. Moreover, generative-AI task exposure should not be equated with physical-robotics exposure: the technologies and affected tasks overlap only partially.

The ILO documents substantial heterogeneity by occupational category and income setting, consistent with variation in access to digitized work, technology, and labour-market structure [74]. How exposure affects actual employment depends upon investment, task complementarity, training, firms' strategies, worker protections, demand responses, and the creation or reorganization of related activities. A historically grounded argument should consequently investigate both the production of new machinery and the reproduction of skills, livelihoods, care, and institutional bargaining capacities. An industrial process can reduce the time required for a narrow task while increasing the number of workers required for system integration, repair, verification, or material provision at another site. The direction and distribution of these changes must be established empirically rather than read directly from task-level technical capability.

The historical-materialist significance extends beyond counting jobs. Productive intelligence is socially formed through education, research, cooperation, and practical work, as discussed in Section 3. The incorporation of some of that intelligence into machinery can change who performs which operations and who controls the productive apparatus. But the reproduction of working and social capacities involves housing, health, education, unpaid care, collective institutions, and access to income. The organization of these conditions mediates the consequences of technical change. A later empirical study should therefore include conditions of work, compensation, skills, employment transitions, and social reproduction as outcomes alongside output, uptime, electricity, and resource efficiency.

## 5.8 Identifying the Material Reproduction Problem

The indicators collected above are mutually relevant without belonging to a single directly comparable statistical series. Industrial robots are counted in installed or operational units; electricity in TWh per year; investment in nominal currency; extracted materials and waste in tonnes per year; refinery concentration in market shares; and labour exposure in shares of employment. Any composite index that combined these magnitudes without explicit normalization, causal theory, and appropriate weights would conceal important assumptions. More fundamentally, an accounting identity about observed material flows does not by itself determine which institutional arrangements produced them or which ecological consequences they generated.

The historical-materialist investigation therefore begins with a specified productive function, a system boundary, an institutional configuration, and a set of counterfactual conditions. For example, a comparison of two systems producing an equivalent quantity of qualified industrial components should

DRAFT

separately document the machinery and workers required, outsourced model and compute services, energy consumption over time, upstream materials, repair and replenishment requirements, equipment retirement, and ecological impacts associated with those flows. Comparisons should hold output quality and relevant operating conditions sufficiently constant, or model the consequences when they differ. The production relation should be described through observable access rights, ownership, decision authority, financing, and the distribution of gains and burdens, rather than attributed to a unit simply because it employs AI.

Three empirically distinguishable questions follow. First, does greater operational autonomy reduce the demand for particular kinds of human intervention while increasing the dependence on specified external components, labour, computing, or energy? Second, does a particular change in automation reduce physical inputs or lifecycle impacts *per comparable unit of output*, and what happens to aggregate impacts when output, prices, or consumption change? Third, do differences in ownership, access and governance alter the set of feasible production policies under the same engineering and ecological constraints? These questions require facility-level panel data, supply-network information, life-cycle and material-flow inventories, reliable operational records, and transparent comparison groups. If the proposed relational representation adds no explanatory or predictive capability beyond competent established methods, that negative result would be part of the research finding.

This empirical baseline therefore warrants a precise conclusion. Contemporary AI-mediated production is developing within an unevenly distributed stock of machinery and infrastructures, measurable electricity and investment demands, concentrated stages of material processing, incomplete recovery systems, and heterogeneous labour arrangements. These conditions establish neither the arrival of fully autonomous productive reproduction nor a predetermined ecological outcome. They identify the material relations that a theory of productive autonomy must explicitly represent and test. Section 6 consequently introduces a relational representation of production dependencies designed to preserve heterogeneous physical units, institutional control relations, boundary choices, and uncertainty, rather than substituting conceptual integration for empirical evidence.

# 6 A Relational Representation of Production Dependencies

The historical and empirical discussion in Sections 3–5 established that an automated production facility depends upon material equipment, energy, knowledge, labour, coordinated organizations, and ecological conditions whose provision may occur beyond its immediate operating boundary. This section develops a representation in which such dependencies can be distinguished, recorded, and compared. Its purpose is more limited than deriving a comprehensive ontology of production: it specifies an investigable set of relations that can connect physical production records to historically variable institutional arrangements, while identifying what additional evidence would be required for causal or normative claims. The key methodological choice is to preserve heterogeneity. A quantity of electricity, an inventory of components, a legal permission, an asserted scientific fact, and a habitat condition cannot be treated as quantities of one interchangeable substance simply because each can affect the feasible continuation of a production process.

## 6.1 The Scope of a Production Relation

A productive dependency is a claim that the occurrence or continued feasibility of one activity is conditioned by a resource, capacity, event, or relation elsewhere in a production system. The claim can concern a physically necessary input, an organizational prerequisite, a probabilistic reliability dependence, a contractual condition, or an empirically inferred ecological interaction. These are distinct relations, even when they involve the same organizations and equipment. In particular, a recorded supply contract does not demonstrate delivery of a physical input; measured delivery does not demonstrate lawful authorization; and the association of an ecological change with a productive activity does not identify a causal

DRAFT

effect without an adequate research design. The representation must permit these distinctions to remain visible during analysis.

There are useful precedents. Leontief's input–output analysis gives a quantitative representation of inter-industry flows, while environmentally extended multiregional systems such as EXIOBASE link production accounts to environmental pressures [112, 173]. Such historical, sector–region accounts cannot by themselves resolve the observed equipment dependencies, supplier identities, or ecological exposures of a particular contemporary plant. Life-cycle assessment requires explicit goal and scope, inventory, impact assessment, and interpretation choices [84]. Formal manufacturing models, including coloured and timed Petri nets, can make process sequencing, resource use, and discrete state changes explicit [189]. The W3C PROV ontology separately identifies entities, activities, agents, and relations of use, generation, and responsibility [196]. The ISO 23247 manufacturing digital-twin framework supplies another relevant architectural comparator: its overview and general principles already organize the description of a manufacturing twin [88]. The proposed generative-relational (GR) description is therefore an exploratory integration layer over these more established methods, not a claim to have originated typed production objects, digital twins, provenance records, or relational modelling. Any claimed added value must be assessed against properly specified implementations of them.

Let a research question specify an observation domain $\Omega$, a system boundary $B$, an observation protocol $\mathcal{O}$, and an ordered collection of analytic intervals $I_0, I_1, \ldots, I_N$. The interval index $n$ is a bookkeeping device that may summarize continuous processes or irregular event histories; it does not impose a universal temporal ontology. In particular, an institution's later reconstruction of an event sequence may differ from what a controller recorded as events occurred. Historical statements about sequence and causality therefore require stated source and dating conventions. Define the provisional relational description

$$\mathscr{P}_{B,\Omega,N} = \big(\mathsf{V}, \mathsf{R}, \tau, \mathsf{F}, \mathsf{C}, \mathsf{H}, \mathcal{O}, \mathcal{E}\big)_{B,\Omega,N}, \tag{1}$$

where $\mathsf{V}$ is a set of typed objects and event records, $\mathsf{R}$ a set of typed relations, $\tau$ assigns types, $\mathsf{F}$ contains physically measured or estimated flows, $\mathsf{C}$ describes capacities and constraints, $\mathsf{H}$ records ordering and dependency histories, and $\mathcal{E}$ contains evidence and provenance claims. This object is a research representation, not an assertion that every represented relation possesses a common metaphysical status or a known causal mechanism.

## 6.2 Typed Entities and Productive Events

The basic types are productive activities $\mathsf{Q}$; actors or organizations $\mathsf{A}$; material inventories and products $\mathsf{M}$; equipment and infrastructures $\mathsf{I}$; datasets, models, specifications, and other representational artifacts $\mathsf{K}$; ecological sites or processes $\mathsf{Z}$; and recorded events $\mathsf{T}$. They can overlap in carefully declared ways: an autonomous laboratory may be treated as an institutionally governed facility at one level and as an interconnected set of instruments, work processes, and services at another. The type assignment $\tau$ is consequently contextual rather than an ontologically privileged partition. Each activity $q \in \mathsf{Q}$ can be represented as a qualified event with input relations, output relations, operative constraints, location, and an actor or control system associated with its execution. Distinguishing *agent associated with an activity* from *agent authorized to decide it* is especially important when automated equipment executes instructions generated within another institution.

Dependencies frequently require several conditions jointly. Manufacturing a qualified component might require a machine, a calibrated tool, a feedstock grade, available electrical power, and an accepted inspection procedure. A binary network edge can represent each condition separately, but a *typed directed hyperedge* more directly describes a productive activity as the joint use of a collection of conditions to generate a collection of outputs. A production requirement can also have alternative satisfiers, such as two certified suppliers, two repair strategies, or several compatible energy sources. Thus, a dependency description includes conjunctions, alternatives, quality constraints, capacity constraints, and possible substitutions. A graph path alone may overstate dependence when substitution is available or understate it when multiple inputs are required simultaneously.

DRAFT

Table 7: Typed relational claims and their respective requirements for empirical interpretation. A common digital record does not make the claims interchangeable.

| Relation type | Illustrative observation or record | What remains necessary for the claimed inference |
|---|---|---|
| Physical transformation | Metered inputs, qualified outputs, measured losses | Conservation or process model, sensor validity, allocation boundary |
| Energetic provision | Time-stamped delivered electricity or fuel | Quality, capacity, upstream generation and conversion losses |
| Technical enablement | Equipment compatibility or service dependency | Tested operating envelope and feasible substitution pathways |
| Information or knowledge use | Dataset, model, specification, or provenance link | Usage evidence, validity, interpretation and intellectual-property scope |
| Institutional control | Title, licence, instruction, approval, contract | Competent source, temporal scope, enforceability and remedy |
| Labour and care contribution | Task logs, training, repairs, maintenance | Work conditions and socially distributed reproductive support |
| Ecological interaction | Exposure, discharge, habitat or biotic observation | A defensible causal design and ecologically appropriate endpoints |

A minimal activity record should declare (i) a unique activity identifier and relevant interval, (ii) input and output types and quantities where physical measurement is appropriate, (iii) enabling equipment and qualified personnel or control policies, (iv) evidentiary sources, and (v) any applicable institutional and ecological constraints. Missing fields are not silently filled by treating all undocumented relations as absent. An unrecorded subcontractor, upstream electricity load, or ecological effect may be consequential even when unavailable to the investigator. Evidence completeness and the completeness of the physical production system are therefore separate properties.

Table 7 organizes the representation into relations that can, in principle, be independently assessed. It also helps distinguish descriptions of productive capacity from evaluations of who benefits, who bears ecological burdens, and who has the legitimate authority to authorize a particular activity. These latter assessments require theories and procedures beyond a typed data structure.

### 6.3 Dimensional Accounting of Material and Energy Flows

Physical quantities require additional structure. Let $x_{k,n}$ denote an inventory of commodity or substance $k$ at the beginning of analytic interval $n$, with unit $[x_k]$ explicitly specified (for example kilograms of a material grade or counted components). Let $u_{j,n} \geq 0$ denote the number of executed cycles of process $j$ within that interval. The coefficient $S_{kj,n}$ gives the net change in inventory $k$ per valid process execution, expressed in units $[x_k]/\text{cycle}_j$. The following relation is a bookkeeping identity *only if* the full set of sources, sinks, and inventory changes has been recorded consistently:

$$x_{k,n+1} - x_{k,n} = \sum_{j \in \mathsf{Q}} S_{kj,n} u_{j,n} + r^{\text{in}}_{k,n} - r^{\text{out}}_{k,n} - \ell_{k,n}, \qquad k \in \mathsf{M}. \tag{2}$$

Here the $r$ terms are recorded boundary transfers and $\ell_{k,n}$ an explicitly defined loss or inventory correction, all measured in $[x_k]$ over the interval. Where a process changes the composition of material, the commodity balance must be connected to appropriately measured material or elemental conservation constraints. Counting one finished component and kilograms of feedstock in the same untyped inventory would render Equation (2) meaningless. Conversion coefficients, tolerances, inventory uncertainty, co-products, scrap and embedded material flows therefore belong in the data model. Numerical process yields should be empirically estimated within stated operating conditions.

Energy accounting likewise requires compatible dimensions. Let $e_{j,n}$ represent energy used per execution of process $j$, measured in megajoules per cycle, $E^{\text{aux}}_n$ auxiliary energy use in megajoules, and

DRAFT

$E_n^{\text{sup}}$ the energy available within the same interval. An operational necessary condition is

$$\sum_{j \in \mathsf{Q}} e_{j,n} u_{j,n} + E_n^{\text{aux}} \;\leq\; E_n^{\text{sup}}, \tag{3}$$

when all three magnitudes cover the same boundary and interval. Equation (3) is not sufficient for operability: instantaneous power limits, peak demand, voltage quality, outage duration, temporal matching, storage, and thermal constraints may invalidate a schedule even when its total energy requirement is covered. Likewise, electricity consumed by a remote model-serving facility must be included when the analytical question concerns extended infrastructure, and its allocation requires a documented rule. These distinctions are consistent with the goal-and-scope discipline of life-cycle accounting [84].

Information is recorded on different terms. A file size can be quantified in bytes, latency in seconds, and a trained model's computational footprint through declared metrics; a worker's practical knowledge or a research programme's epistemic warrant cannot in general be reconstructed from these quantities. The expression "information flow" therefore refers to documented use or transmission events unless an operational information measure is specifically defined. Economic values, measured in currencies at defined dates and under defined accounting rules, remain separate again. Price-weighted aggregation of distinct physical stocks may serve a declared economic question but cannot establish material substitution or normative commensurability by arithmetic alone.

## 6.4 Productive Feasibility and Institutional Authority

Engineering feasibility depends on capacities, timing, compatibility and resource constraints. Institutional authority can restrict or enable which technologically possible operations are admissible. Define $\mathcal{U}_n^{\text{phys}}(x_n)$ as the set of process schedules satisfying the specified physical and technical constraints, and let $\mathcal{U}_n^{\text{inst}}(a_n, h_n)$ contain schedules admitted by the applicable, *verified* institutional conditions for organizational configuration $a_n$ and recorded history $h_n$. For a declared research question, the feasible set may be represented as

$$\mathcal{U}_n^{\text{feas}} = \mathcal{U}_n^{\text{phys}}(x_n) \cap \mathcal{U}_n^{\text{inst}}(a_n, h_n) \cap \mathcal{U}_n^{\text{eco}}(z_n; \mathcal{C}), \tag{4}$$

where $\mathcal{U}_n^{\text{eco}}$ represents only those ecological constraints whose content, evidence and governing status have been explicitly specified. Ecological admissibility may include legal thresholds, research-protocol limits, or scientifically motivated scenario constraints; these sources should not be conflated. The intersection records compliance with the *declared* restrictions, not the provenance or legitimacy of each restriction. A legal permission belongs to the institutional account of authorization even when its ecological purpose motivates a constraint in $\mathcal{U}_n^{\text{eco}}$; any such overlap must be identified rather than counted as independent scientific evidence. A mathematically convenient limit on habitat exposure is not, by itself, a democratically authorized ecological standard or evidence that a particular threshold protects all relevant organisms.

An access relation can be coded as $R_{a,i,n}^{\text{access}} \in \{0, 1, ?\}$, where the final symbol indicates unresolved evidence rather than refusal. The coding needs to reference the applicable licence, contract, permit, technical authentication condition or other instrument; the identity of its competent issuer; and any geographical, temporal, operational or revocation conditions. Institutional records may conflict or be contested. For such cases, multiple sourced assertions and their status should be retained rather than collapsing them into a fictitious single ground-truth flag. In economic analysis, control over an indispensable production relation may shape an actor's bargaining position, but a quantified measure of power would additionally require specified feasible alternatives, transaction costs, substitution constraints and institutional responses.

**Proposition 1 (Dependence of feasible production on institutional conditions).** Consider two otherwise identical descriptions with the same physical and ecological feasible sets and with institutionally admissible sets $\mathcal{U}_1^{\text{inst}} \subseteq \mathcal{U}_2^{\text{inst}}$. Their resulting feasible sets satisfy $\mathcal{U}_1^{\text{feas}} \subseteq \mathcal{U}_2^{\text{feas}}$.

DRAFT

*Proof.* Each feasible set is the intersection of the common physical and ecological sets with its respective institutional set. Intersection preserves set inclusion. □

This elementary result is a statement about the representation, not an empirical theory that wider access rights always increase realized output, improve working conditions, or diminish ecological burdens. The result also depends on what counts as an admissible institutional condition. Changing governance may alter investment, incentives, coordination, and the physical or ecological feasible sets themselves. Those effects cannot be derived by holding them fixed, and they remain matters for comparative historical and empirical investigation.

### 6.5 Ecological Dynamics Within the Productive Boundary

Physical production and ecological continuation are linked without sharing an automatically commensurable value metric. Let $z_n$ represent an evidence-dependent vector of ecological observations or inferred state variables, such as substrate condition, water use, contaminant concentration, habitat structure, or taxon-specific indicators. A conditional dynamic description may be written

$$z_{n+1} = f_{\theta,n}(z_n, u_n, d_n, \varepsilon_n), \tag{5}$$

where $d_n$ represents other disturbances, $\theta$ specifies a model family or calibrated parameters, and $\varepsilon_n$ denotes residual influences under a stated probabilistic or uncertainty model. Equation (5) has no ecological predictive validity before relevant variables, spatial support, observation errors, parameter identification, and external stressors have been established. A technologically efficient operation can generate indirect changes beyond the facility's boundary, and the representation should retain the possibility of cumulative effects and delayed or nonlinear responses.

The system boundary plays different roles in physical inventory accounting, ecological interpretation, and governance. An inventory boundary determines the activities whose resource flows are counted. An ecological influence boundary may be wider because upstream resource extraction, emissions, hydrological interactions, and supply chains extend beyond the operating site. A jurisdictional or contractual boundary can be different again. A study reporting lower electricity consumption within a factory gate cannot infer lower lifecycle electricity demand without examining upstream computation and outsourced processing. A study observing no short-term ecological change around a facility cannot infer the preservation of regional ecological reproduction without suitable ecological coverage. Such distinctions build on established lifecycle and industrial-ecological accounting and should be made explicit before proposing any new GR measure [84, 173].

### 6.6 A Changing Network of Reproductive Dependencies

Historical materialism concerns evolving productive arrangements, including changes in who controls equipment, whose labour reproduces productive capacities, and where materials and knowledge are obtained. A static graph of observed supply relationships can serve a snapshot, but it cannot alone represent supply substitution, equipment aging, institutional changes, or the emergence and disappearance of production activities. Denote the typed dependency structure during interval $n$ by $\mathcal{G}_n$. A provisional transition description is

$$\mathcal{G}_{n+1} = \Phi_{\psi,n}\big(\mathcal{G}_n, u_n, a_n, \omega_n\big), \tag{6}$$

where $u_n$ contains recorded production or investment actions, $a_n$ relevant organizational changes, and $\omega_n$ unmeasured disturbances or exogenous changes relative to the chosen boundary. The transition operator $\Phi_{\psi,n}$ is a modeling hypothesis: it is neither a historical law nor evidence that technological development follows a predetermined direction. Its parameters and permissible transitions would need documentary, economic and engineering support.

A reproductive dependency differs from a dependency needed to execute one process cycle. A robot may produce an acceptable part whenever a calibrated actuator, control system and power supply are available, while restoring those capacities after wear or failure requires further processes: manufacture

DRAFT

of replacement bearings, preparation of qualified calibration equipment, renewal of electronics, training of maintainers, and continued supply of electricity. These processes form another layer of the production network, which in turn depends on other inputs and organizations. In Sections 7 and 8 the distinction between operation and reproduction will be formalized with a declared time horizon and a bounded external-support set, rather than assuming that recursive dependence can be terminated by choosing a convenient accounting boundary.

The representation must distinguish a pathway that is *historically realized* from one that is merely technically feasible. A potential substitute supplier is not necessarily accessible at a viable price, delivery interval or quality level. An open-source design does not guarantee access to manufacturing tools. Conversely, a material bottleneck apparently concentrated at one node may have multiple practical alternatives once changes in design, repair or organization are admitted. Feasible dependencies should therefore be evaluated under specific, time-indexed capacities and institutional conditions, with historical changes treated as empirical questions.

## 6.7 Observation, Provenance, and the Limits of Identification

A production graph is ordinarily assembled from incomplete observations. Equipment logs, enterprise databases, environmental sensors, invoices, contracts, interviews and scientific publications capture different aspects of productive relations. The evidence associated with each assertion should record the source, event time or interval, method of measurement or interpretation, transformation history, and uncertainty where available. W3C PROV-O provides a standardized provenance ontology for entities, activities, agents, usage, generation, and attribution; the proposed representation should interoperate with it where possible [196]. Provenance records explain the basis and history of a claim, but they do not certify that a sensor is calibrated, that an economic attribution is causally justified, or that a purported authorization is legally effective.

If $X_n$ denotes the latent set of conditions of interest and $Y_n$ the observations under measurement protocol $\mathcal{O}_n$, an operational abstraction is

$$Y_n \sim K_{\mathcal{O}_n}\big(\cdot \mid X_n, \pi_n\big), \tag{7}$$

where $\pi_n$ denotes the observation or inspection policy and $K_{\mathcal{O}_n}$ its conditional observation model. In some applications, $X_n$ may be a vector of measurable physical quantities; in others it is a structured collection of inferred relations. Changing the inspection policy can change the records produced even when the underlying productive conditions remain comparable. This includes selective quality sampling, monitoring concentrated at first-tier suppliers, and omission of informal maintenance or outsourced compute. The model is useful only if selection mechanisms and measurement errors can be investigated with sufficient data.

A basic identification limit follows: if two distinct feasible dependency structures induce the same distribution of all observations permitted by a chosen protocol, that protocol alone cannot distinguish the structures. For example, a facility's operating log may show the same output and uptime whether the critical replacement part is acquired through a diversified supplier network or a single inaccessible upstream producer. The difference becomes observable only through suitable supplier and repair evidence or an informative intervention. Additional graph complexity cannot solve a lack of identifying observations. Negative results should be expected whenever the proposed GR recording adds descriptive fields but leaves such evidence gaps unchanged.

## 6.8 A Bounded Illustrative Production Case

Consider a hypothetical cell producing one accepted precision component per recorded activity cycle. Assume for illustration that it uses $0.12$ kilograms of one specified input material grade and yields a $0.10$ kilogram qualified component and $0.02$ kilograms of recoverable scrap, with no other measured material output for this simplified account. Let the same cycle require $2.0$ kWh of on-site electricity and an

DRAFT

allocable $0.4$ kWh of external computing electricity. These numbers are *constructed for exposition*; they are not measurements of an existing facility. Under the simplified boundary, a physical mass account closes at $0.12 = 0.10 + 0.02$ kg. Electricity amounts total $2.4$ kWh only under a compatible accounting boundary and a documented allocation of remote computation. The material flow does not close the energy or ecological account, and recovering $0.02$ kg of scrap does not establish the eventual recovery yield or the environmental burden of recycling.

Suppose the cell receives materials through either of two nominally compatible suppliers and relies on licensed inspection software. A physical network would record both possible material sources and the inspection activity; a historical realization record would identify which supplier actually delivered the input and whether the software was available during the interval. A control record would distinguish permission to run the software from ownership of the equipment or an intellectual-property claim to the design. If the inspection licence expires or a sole certification service becomes unavailable, the feasible production set may contract even though the inventory and energy accounts remain unchanged. Conversely, a validated open replacement for the inspection software could enlarge feasible production options without changing the immediate resource balance. Neither case implies improved ecological outcomes without data on investment, quality, workloads and indirect demands.

The illustrative cell makes an incremental-value test possible. Investigators should first define a common held-out task, such as predicting qualified-output interruption following a repair-part failure or reconstructing the authorized route by which a supplier was changed. A competent process/reliability model, an established production-network or life-cycle workflow, and a GR-augmented workflow should initially receive the same underlying records and uncertainty descriptions. Their errors, calibration, abstentions, analyst time, and maintenance costs can then be compared. A second, separately reported comparison may allow new institutional-event records to be collected for the GR workflow, but must include the additional acquisition cost and an analogous opportunity to enrich competing methods. Otherwise a claimed representational gain could merely be an advantage in data access. Validation should use subsequent events, externally adjudicated incidents, or documented interventions rather than in-sample narrative fit. No such evaluation is reported in this manuscript.

### 6.9 Conditions for an Empirically Useful Representation

The proposal is consequently an interface for joint questions whose evidence crosses established disciplinary databases, not an additional physical law of production. Its value depends on whether typed alternatives and conjunctions of requirements, temporal realizations, governing rights, and source uncertainty improve a declared prediction or decision at a measurable cost. Where a Petri net, input–output table, standard life-cycle inventory, or conventional event database performs the task as well, that method supplies the appropriate analysis. Where the added institutional or historical fields improve performance, the finding identifies a task-specific benefit rather than validating the GR framework as a general ontology. A further comparator is a manufacturing digital twin conforming to a declared implementation of the ISO 23247 framework, equipped with the same process-event, provenance, maintenance and authorization records as the GR candidate [88, 196]. If a well-specified twin or temporal production graph can express the required conjunctive and alternative dependencies and attain equivalent performance, the proposed GR syntax has not established independent predictive value. An advantage claimed merely because an existing comparison model was denied relevant records is a data-availability effect, not a validated representational result. In the present paper, cross-domain integration is a programme-level research hypothesis; neither formal construction alone nor the coexistence of physical and institutional relations establishes scientific novelty.

An independent methodological contribution would require an advantage associated with a *specified operation*, not simply a more inclusive vocabulary. For each prespecified task $q$, an evaluator should define a verified outcome $y_q$, a task-specific loss $L_q$ (including a rule for abstentions), an admissible record set $D_q$, and a shared resource budget. If $\widehat{y}_{q,B}$ and $\widehat{y}_{q,G}$ are the results of a competent baseline and

DRAFT

a GR-augmented implementation using $D_q$, respectively, one prospective comparison is

$$\delta_q = L_q(\widehat{y}_{q,B}, y_q) - L_q(\widehat{y}_{q,G}, y_q). \tag{8}$$

Positive $\delta_q$ indicates an improvement only on that task and loss definition; it is not itself proof that the GR representation is more expressive or that its causal account is correct. Where the loss is measured in hours, dollars, or classification errors, the resulting contrast has that task's units and cannot be pooled directly with differences measured in other units. The evaluator must report uncertainty over genuinely independent held-out cases, leakage controls, computing and analyst costs, and whether the same improvement is obtained by extending the baseline with the identical typed information.

Three outcomes therefore need to be distinguished. *Representational necessity* would require showing that a precisely specified query or invariant cannot be represented by a defined comparator within explicit modelling restrictions; a mere failure to implement it is insufficient. *Task-level empirical utility* requires a repeatable advantage against capable, equally informed alternatives at a declared cost. *Institutional usability* concerns whether the resulting evidence can be interpreted, contested, corrected, and maintained by authorized participants; that assessment is procedural and cannot be inferred from prediction accuracy. None of these tests has yet been carried out here, and failure of one does not automatically decide the others.

This section establishes the representational foundation for the paper's next arguments. Operational autonomy can now be assessed over specified activities and boundaries, with explicit attention to the support necessary to execute and restore a process. Recursive reproduction can be studied through dependencies on maintenance, component production, scientific knowledge, labour and ecological processes without treating those heterogeneous conditions as commensurable. The next section accordingly asks how autonomy should be defined within a given production boundary and evaluation horizon, and which forms of continued external support remain compatible with that definition.

# 7 Defining Autonomy Within Production Boundaries

The relational representation established in Section 6 distinguishes physical inputs, organizational prerequisites, institutional permissions, ecological conditions, and the evidence through which they are known. This section uses that representation to specify what it could mean for an AI-mediated production system to be autonomous. Its immediate purpose is conceptual and formal: to define autonomy relative to declared productive tasks, system boundaries, disturbances, external support, and assessment horizons. Its broader purpose is historical-materialist: to explain why reduced direct human intervention in a particular task does not by itself establish independence from the collective labour, accumulated infrastructures, and material conditions that make the task possible. The section develops a conditional account of an *operational–reproductive autonomy gap*, supplies a deliberately elementary proposition, and identifies the measurements needed to determine whether such a gap occurs in actual production systems.

## 7.1 The Relation between Task Automation and Productive Autonomy

The allocation of decisions and actions between humans and machinery is a familiar question in automation research. Parasuraman, Sheridan, and Wickens distinguish information acquisition, information analysis, decision and action selection, and action implementation, with each function admitting different degrees of automation [149]. Their scheme does not imply that an entire productive organization must receive one automation level. A robot may execute a toolpath without an operator while process specification, fault diagnosis, quality adjudication, maintenance, and capital allocation are undertaken by other agents. Research on human–autonomy teaming likewise treats interdependence as a constitutive feature of effective human–machine systems, rather than equating useful autonomy with the absence of cooperation [137, 106]. These precedents motivate a task- and responsibility-specific definition in place of an unqualified assertion that a factory or laboratory is autonomous.

DRAFT

Manufacturing architectures make the difference concrete. The ISA-95 family specifies interfaces among the physical process, sensing and control, operations management, and enterprise-level functions [99]. The existence of a machine-level controller does not imply that inventory replenishment, maintenance organization, production scheduling, or business authorization occurs within that controller. The standard is cited here as an organizational and information-interface precedent, not as a standard that certifies autonomous reproduction. The object to be assessed is accordingly a *specified system performing specified activities under specified conditions*; the word "autonomous" without those qualifiers is too ambiguous to carry the paper's empirical claims.

In historical terms, these specifications matter because automation often reorganizes rather than simply eliminates coordination. Work formerly performed at the machine may reappear in data preparation, remote engineering, infrastructure maintenance, materials supply, or verification. An autonomy claim concerning on-site execution can be true while its interpretation as independence from socially organized production is false. That difference is analytically important whether the facility is privately owned, publicly organized, cooperatively managed, or maintained as a commons-oriented infrastructure.

## 7.2 Boundaries, Activities, and Assessment Conditions

Let $B$ designate a declared production boundary, including its internal facilities, stocks, agents, equipment, and control interfaces. Boundary membership is indexed to an inquiry and may differ from the location of legal ownership or accounting consolidation. Let $\mathcal{T}$ be a finite set of required productive tasks with explicit output quantity, quality, throughput, safety, and recording conditions. Let $H = (I_0, \dots, I_N)$ identify assessment intervals, $\Xi_{B,H}$ the admissible disturbance histories, and $\Pi_{B,\mathcal{T}}$ the policies available under the declared technological and institutional arrangements. A policy is non-anticipating: decisions during $I_n$ may depend on observations then available and earlier recorded histories, without assuming access to future disturbances. This notation follows the interval convention introduced in Equation (1); it is a model of observable ordering, not a claim that time has a universal discrete constitution.

The tasks and conditions define an operational-success predicate $\mathsf{Op}_{B,\mathcal{T},H}(\gamma) \in \{0,1\}$ for a realized trajectory $\gamma$. Success depends on the stated output specification and the requirements judged necessary for the investigation. Failure modes can include incorrect products, missed throughput commitments, hazardous operation, or breakdown. A successful run is thus a claim about performance under a declared testing procedure, rather than a judgement about who has authority to command production or whether the wider production system is ecologically acceptable. Here *operational* refers to demonstrated execution within the task and observation boundary; it is not a synonym for ownership independence, full material self-sufficiency, or freedom from all human oversight.

External support crossing $\partial B$ is represented by a *typed* vector, not one aggregate number. For a specified channel set $\mathcal{R}_B$, write

$$u_{B,n} = \big(u_{r,n}\big)_{r \in \mathcal{R}_B}, \qquad 0 \le u_{r,n} \le b_{r,n}, \qquad n = 0, \dots, N, \tag{9}$$

where $b_{r,n}$ is an available external-support budget and the unit of $u_{r,n}$ depends on its channel. Channels might quantify manual intervention in person-hours, incoming replacement components in units, externally supplied electricity in kWh, model-serving access in compute-hours, or material replenishment in kg of a stated substance. Some provisions, including legal authorization, expert judgement, habitat status, and obligations of care, cannot responsibly be reduced to additive quantities of support. They remain separate admissibility predicates. Importing ambient solar radiation, drawing water, or extracting materials can be physically essential even where no commercial invoice or human delivery is recorded. A zero value for a measured supplier channel is consequently not evidence of physical independence from the environment.

For an admissible policy $\pi$, denote the support schedule used under disturbance $\xi$ by $u^{\pi}(\xi)$ and the resulting trajectory by $\gamma^{\pi}(\xi)$. Support budgets are upper bounds on what may be drawn, not instructions to consume all available support. An idealized robust operational-feasibility set is

$$\mathcal{B}^{\mathrm{op}}_{B,\mathcal{T},H} = \big\{ b : \exists \pi \in \Pi_{B,\mathcal{T}},\ \forall \xi \in \Xi_{B,H},\ u^{\pi}(\xi) \preceq b \ \wedge\ \mathsf{Op}_{B,\mathcal{T},H}\big(\gamma^{\pi}(\xi)\big) = 1 \big\}, \tag{10}$$

Table 8: Distinct autonomy claims and the external supports they require researchers to investigate. These are assessment categories, not empirically calibrated ranks.

| Capacity | Claim being assessed | Illustrative external dependency |
|---|---|---|
| Operational | Runs specified tasks to declared quality and safety thresholds | Intervention on faults; consumable replenishment |
| Planning and adaptive | Selects and revises production policies within its allowed action space | Human-approved objectives; external models and feedback |
| Diagnostic | Detects anomalies and chooses bounded responses | Remote expert diagnosis; reference calibrations |
| Maintenance | Restores equipment to a defined usable condition | Spare parts, specialist tooling, external repair labour |
| Energetic | Maintains required energy quality, power and availability | Grid imports, storage replacement, upstream fuel or energy equipment |
| Material | Secures suitable process inputs and manages losses | Mining, refining, purchased feedstocks, transport |
| Reproductive | Restores or renews required productive capacities over the horizon | Component manufacture, skills, fabrication and other upstream capacities |

where $\preceq$ is the componentwise order over identically typed support channels and assessment intervals. The feasibility of actual deliveries, energy availability, permitted actions, and other conditions is included in the admissible policy and disturbance models. Because reliable knowledge of all disturbance histories is rarely available, Equation (10) is a formal benchmark; practical work may instead report the empirical frequency of meeting task specifications, with confidence intervals and a documented disturbance distribution. Robust and probabilistic claims must be named separately.

This definition also reveals an important consequence of system-boundary choice. Expanding $B$ to include a repair workshop can convert an imported repair service into an internal operation, while introducing the workshop's own needs for labour, tools, materials, and energy. Such a change alters the support channels, policies, and accounting perimeter. Autonomy measures associated with two different boundaries generally cannot be compared merely by subtracting the number of external transactions.

### 7.3 The Distinct Capacities Involved in Productive Autonomy

Autonomy is better represented by a structured assessment of functions and dependencies than by a universal scalar score. Table 8 distinguishes dimensions that can be assigned different tests for the same installation. Each dimension is conditional on its specified task, boundary, horizon, disturbance range, and allowable human or external support. The distinctions are analytic and can overlap in practice.

A laboratory robot can have high operational autonomy within a repeated synthesis protocol and lower autonomy in fault recovery or preparation of new experiments. A solar-powered installation can supply much of its own operating electricity but depend on geographically concentrated producers of electronics and batteries. A factory can maintain a stock of replacement parts without being able to reproduce the productive capacity of the facilities that manufacture those parts. None of these observations provides a general ranking of the technologies; they identify different activities and conditions whose realization would require different evidence.

Autonomy also has an institutional dimension, but institutional independence must be handled carefully. A regulator's approval, a community's participation rights, or a worker's authority to halt an unsafe operation is not a "support deficit" to be engineered away. The ability to execute a task without seeking immediate human permission is an operational characteristic; the legitimacy of that delegation is a distinct claim about authority and responsibility. Treating procedural constraints as losses in a technological autonomy score would silently substitute an optimization criterion for an institutional judgement. The formalism therefore records authorized policy sets separately from the support budget in Equation (10).

DRAFT

## 7.4 Reproductive Autonomy and the Continuation Requirement

A successful run can consume inventories, reduce remaining equipment life, degrade calibration, or depend on unrecorded expertise. These changes may leave the facility less capable of subsequent operation, even if its immediate output meets specification. To separate execution from continuation, let $x_N$ be the measured or estimated terminal productive state. Let $K_{B,\mathcal{T},H}$ be a *predeclared continuation-readiness set*: a collection of states with adequate component health, available inventories, actionable knowledge, verified data or model access, and maintenance arrangements for a specified subsequent production requirement. Membership in this set may be tested through engineering thresholds or a further controlled operating campaign. It is not claimed to be a complete ecological or economic viability kernel.

Define reproductive success by the conjunction

$$\mathsf{Rep}_{B,\mathcal{T},H}(\gamma) = 1 \quad \Longleftrightarrow \quad \mathsf{Op}_{B,\mathcal{T},H}(\gamma) = 1 \quad \wedge \quad x_N(\gamma) \in K_{B,\mathcal{T},H}, \tag{11}$$

and let its associated robust support-feasibility set be

$$\mathcal{B}^{\mathrm{rep}}_{B,\mathcal{T},H} = \left\{ b : \exists \pi \in \Pi_{B,\mathcal{T}},\ \forall \xi \in \Xi_{B,H},\ u^{\pi}(\xi) \preceq b \ \wedge\ \mathsf{Rep}_{B,\mathcal{T},H}\big(\gamma^{\pi}(\xi)\big) = 1 \right\}. \tag{12}$$

**Proposition 1** (Conditional inclusion of feasible support budgets)**.** *For fixed $B$, $\mathcal{T}$, $H$, $\Xi_{B,H}$, $\Pi_{B,\mathcal{T}}$, support-channel definitions, and success predicates as specified above,*

$$\mathcal{B}^{\mathrm{rep}}_{B,\mathcal{T},H} \subseteq \mathcal{B}^{\mathrm{op}}_{B,\mathcal{T},H}. \tag{13}$$

*Proof.* Take any $b \in \mathcal{B}^{\mathrm{rep}}_{B,\mathcal{T},H}$. By Equation (12), a policy $\pi$ exists that meets the external-support bound for every admissible disturbance and produces a trajectory satisfying $\mathsf{Rep} = 1$. Equation (11) includes $\mathsf{Op} = 1$ in that condition. The same policy and disturbance qualification therefore establish membership in $\mathcal{B}^{\mathrm{op}}_{B,\mathcal{T},H}$ by Equation (10). □

Proposition 1 is a consequence of the definitions, not an empirical discovery that existing AI production systems necessarily exhibit a gap. Under fixed specifications, a budget in $\mathcal{B}^{\mathrm{op}} \setminus \mathcal{B}^{\mathrm{rep}}$ witnesses *operational feasibility without verified reproductive feasibility* at that support limit. Strict inclusion is a research hypothesis requiring evidence about deterioration, replacement requirements, supply capability, and available alternative policies. Empty feasibility sets, poorly specified continuation thresholds, or changes in the disturbance range can also affect the comparison. The proposition should not be read as claiming that every production installation fails reproductive testing. Consider the limiting case of a fully stocked, qualified production cell whose components and readiness thresholds remain adequate throughout a finite registered horizon under all modeled disturbances. If its operational and reproductive conditions coincide within that horizon, its two support-feasibility sets can be equal. Thus, strictness requires a documented separating support budget and an actual continuation requirement; a longer observation horizon may reveal a gap, while a short or undemanding terminal audit may conceal one. This constructed case challenges the universal interpretation of an operational–reproductive gap while preserving the conditional inclusion proposition.

A single scalar "autonomy percentage" can obscure this ordering. Because imported electricity, replacement parts, remote labour, and upstream fabrication have different units and substitutability conditions, there may be several incomparable minimally sufficient support profiles. If an index is desired for a particular policy purpose, its weights, normalizations, and value judgements must be declared and subjected to sensitivity analysis. The physically meaningful result is initially a family of feasible support sets and identified dependencies, not an unexplained weighted total.

## 7.5 A Constructed Illustration of the Operational–Reproductive Gap

Consider a hypothetical automated precision-processing cell that must complete two qualified production cycles during the declared assessment horizon. Each cycle consumes one replaceable process cartridge. The cell begins with exactly two cartridges, has functioning control equipment and sufficient energy, and

DRAFT

Table 9: Constructed two-cycle cartridge example. Operational and continuation tests are fixed while the permitted external replacement support changes; all values are illustrative.

| **Criterion** | **Zero cartridge imports** | **Two qualified cartridge imports** |
|---|---|---|
| Qualified cycles completed | 2 of 2 | 2 of 2 |
| Terminal qualified cartridges | 0 | 2 |
| Operating target met | Yes | Yes |
| Continuation inventory threshold met | No | Yes |
| Warranted conclusion | Bounded operational execution | Bounded execution with external replenishment |

requires no human action during ordinary execution. Suppose that successful continuation into the next two-cycle campaign is operationally defined to require two qualified cartridges available at the end of the assessment period, together with usable machinery. No internal cartridge-fabrication capability exists. All numbers in this example are chosen solely to demonstrate the definitions; they describe no measured factory.

With no external cartridge delivery, a controller can execute both cycles, producing two qualified units under the stipulated absence of other failures. It therefore satisfies the operational specification but ends with zero cartridges. It fails the continuation-readiness test. If an authorized and physically feasible external delivery of two qualified cartridges is permitted before the terminal audit, the same installation may satisfy both conditions. The support difference is measured in cartridge units; it cannot be treated as a direct exchange rate against kilograms of raw material, person-hours, or the energy embodied in cartridge manufacturing. If an internal cartridge recycler were introduced, the result would also depend on actual recovery yields, quality losses, equipment wear, and its new dependencies.

This example reveals two distinct limitations of an unqualified autonomy claim. First, uninterrupted task execution can draw down the conditions for future execution. Second, restoring continuation by importing cartridges may solve the facility-level problem while maintaining dependence on external productive networks. Neither the absence of operator intervention nor the feasibility of replacement establishes the social or ecological acceptability of cartridge production. A comprehensive evaluation must add the supply chain's physical burdens, labour relations, access conditions, and uncertainty, as proposed in later sections.

### 7.6 Boundary Sensitivity and the Redistribution of Dependencies

The dependence on $B$ is analytically consequential. A facility-level record might classify remote maintenance, replacement electronics, model updates, and utility electricity as external supports. A corporation-wide boundary might classify some of the same supports as internal transfers. A regional or international input–output boundary could reveal upstream mining, refining, fabrication, transport, or labour dependencies that neither facility record displays. Each boundary is legitimate for some questions and insufficient for others. Enlarging a boundary can change what is called external support while leaving the physical process and its environmental consequences substantially unchanged.

This observation can be stated as a boundary-comparison warning. Given $B_1 \subset B_2$, neither

$$\mathcal{B}^{\mathrm{op}}_{B_1,\mathcal{T},H} \quad \text{nor} \quad \mathcal{B}^{\mathrm{op}}_{B_2,\mathcal{T},H} \tag{14}$$

can generally be ordered without a specified translation of support channels, inventories, policies, disturbance sets, and responsibilities. Including a supplier within $B_2$ may reduce counted imports of finished components but increase counted requirements for raw materials and specialist machinery. The same holds for reproductive feasibility. Because the two budget spaces may have different dimensions and physical units, a naive scalar comparison is not even mathematically well-typed. The analysis should instead provide a mapping of boundary changes, with matched physical accounting where possible.

DRAFT

This is particularly important in historical materialism. The capacity of a factory to automate local tasks is formed through relations extending beyond its gates: scientific research, training, purchased means of production, global extraction and transport, ownership, and social reproduction of workers and engineers. A boundary choice that excludes those relations can be useful for a local control problem while remaining inadequate to explain the historical reproduction of the productive forces. Conversely, tracing each dependency indefinitely offers no practicable finite measurement protocol. The paper therefore favors nested, declared research boundaries and explicitly reported residual dependencies over claims to have captured an entire totality.

A change in ownership or access can also alter operational feasibility without immediate alteration of machine hardware. Under the typed description of Section 6, revocation of a software licence, interruption of a supplier contract, a safety stop, or lawful withholding of authorization may modify the admissible policy set. Such changes are not reducible to component reliability. Their empirical investigation requires contracts, policy records, governance procedures, and evidence of actual effects on operation. A normative assessment must separately ask whether each form of intervention is legitimate, proportionate, and contestable.

## 7.7 Ecological Conditions and the Limits of Technical Continuity

A reproductive test based only on terminal inventories and component condition remains incomplete for the paper's ecological purpose. *Reproductive autonomy* is assessed relative to a defined horizon and support boundary; it asks how the enabling capacities can be maintained or replenished, not whether a productive system can exist outside physical and ecological exchanges. The system could meet those thresholds while drawing excessive water, accelerating material extraction, or causing damage that burdens external communities or ecological systems. Accordingly, the admissible policy set may be intersected with a declared ecological-feasibility set $\Pi^{\mathrm{eco}}_{B,H}$, consistent with the distinction between technological, institutional, and ecological feasibility already expressed in Equation (4). Researchers must state which ecological indicators are included, how exposure is measured, what system boundaries are used for upstream burdens, and which uncertainties remain unresolved.

This ecological restriction cannot be inferred simply from the presence of renewable electricity, local recycling, or automated maintenance. Every pathway depends on physical energy conversion, materials and infrastructures. Recovering more material can lower some forms of extraction while increasing other processing demands; extended equipment lifetime can reduce replacement production while changing maintenance energy or material requirements. Any net benefit requires context-specific life-cycle assessment and counterfactual comparison rather than a definitional association between autonomy and sustainability [84, 173].

The political-economic implications are equally bounded. Reduced direct labour hours within $B$ neither measures the total living labour involved in the expanded supply system nor specifies the distribution of income, authority, occupational risks, or technological gains. Reproductive capacity may be supported by publicly funded science, care activities, and collective maintenance whose contributions do not appear in a narrow facility ledger. The historical question is therefore how technical autonomy is produced and governed through material and social dependencies, and how different institutional arrangements make different continuations feasible. Ethical acceptability and ecological justice remain questions requiring evidence and reasoned institutional judgement beyond the autonomy set calculations.

## 7.8 Measuring Autonomy in Comparative Production Research

The formal definitions provide a design for empirical investigation rather than ready-made estimates. A comparative study would first specify a reproducible task campaign, declared safety and output thresholds, a clear facility boundary, and at least two assessment horizons. It would then record the occurrence, origin, and causal role of external supports: operator and engineer interventions, replacement components, remote computing, calibration services, grid imports, software dependencies, maintenance processes, and material supply. Event provenance should distinguish planned assistance from unplanned

DRAFT

failure recovery, and measurements from supplier reports or model-based allocations. The proposed GR representation can preserve those distinctions, but conventional maintenance logs, enterprise integration records, life-cycle inventory, and provenance systems provide necessary comparison baselines [99, 196].

The protocol should also carry out fault and interruption scenarios calibrated to realistic risks. A nominal uninterrupted demonstration can be compared with a campaign that includes documented network interruptions, delayed components, equipment faults, or reconfiguration requirements. Researchers can report output quality, task completion, downtime, restoration duration, external person-hours, qualified spare parts consumed, and energy delivered, each in its own unit. Where disturbances are sampled rather than enumerated, uncertainty intervals and sampling assumptions should accompany reliability claims. The continuation-readiness test should be specified *before* the campaign, with independently checkable inventories and equipment-condition measurements. A further run may be needed to establish whether the terminal state really sustains the nominated subsequent task set.

A useful experimental design would compare a facility under three policies: a fixed automation policy, an adaptive policy optimized for nominal task completion, and an adaptive policy explicitly constrained by restoration of productive capacities. All three should face matched workloads and comparable support availability. The analysis would investigate whether increased task-level automation moves the system closer to or farther from its continuation-readiness set, and whether any apparent gap disappears when maintenance is improved or inventory buffers are added. These competing explanations are important: a gap might be caused by avoidable scheduling choices or measurement omissions rather than a structural feature of AI technology. Ecological indicators and external supply-chain impacts should be measured separately, not folded into task success.

The proposed approach is falsifiable. If repeated campaigns find that operational improvements systematically preserve continuation readiness, diversify effective suppliers, or reduce system-wide ecological burdens under appropriately matched conditions, a strong hypothesis of autonomy-induced reproductive fragility would be contradicted for those systems. If the GR record does not improve reconstruction of critical dependencies or anticipation of failures beyond competent conventional records, its asserted empirical advantage should likewise be rejected. Such outcomes would refine the research programme rather than invalidate the legitimacy of studying autonomous production.

### 7.9 Implications for the Reproduction of Productive Capacity

Autonomy in production can now be expressed as a bounded family of task-performance and support-feasibility claims. The operational–reproductive distinction makes visible a possible divergence between executing a productive activity and restoring its conditions of continuation. The conditional inclusion result establishes how the two tests relate by definition, while the existence, magnitude, causes, and consequences of an actual gap remain empirical matters. A facility's ability to operate without immediate human intervention may be genuine and significant even when it continues to depend on other workers, institutions, material suppliers, and ecological systems.

This account provides a transition to Section 8, which examines the deeper problem of recursively reproducing infrastructure and the limits of closure across productive levels. The next question concerns the conditions under which a system that replenishes its own immediate consumables can also maintain the equipment, knowledge, and productive relations required to replenish those consumables in subsequent cycles.

## 8 The Recursive Reproduction of Production Infrastructure

Section 7 distinguished the successful execution of a specified production campaign from restoration of the capacity to execute another campaign. This section investigates the deeper dependency structure implicit in that distinction. Its object is the *recursive reproduction of production infrastructure*: the processes through which machinery, tools, energy provision, feedstocks, information infrastructures, and the capabilities needed to maintain them are replenished through further productive activities. The argument

DRAFT

develops a boundary- and horizon-specific account of such reproduction, connects the historical problem to established interindustry and industrial-ecological analyses, and identifies the assumptions under which a quantitative upstream-dependency calculation is legitimate. The resulting propositions characterize an accounting model, while claims concerning actual infrastructural autonomy remain questions for empirical investigation.

## 8.1 Reproduction beyond the Immediate Productive Cycle

A production cycle consumes more than the materials incorporated into its final outputs. Tools wear, calibration drifts, filters and lubricants require replacement, control systems depend upon updates and qualified data, and productive facilities draw on electric-power and communication networks. Some of these requirements are intermittent, and some arise only following failures. A facility can therefore complete its assigned production while reducing the stock of capabilities required for future cycles. The terminal continuation-readiness test introduced in Section 7 records one aspect of this difference. Yet replenishment of a depleted component poses a further question: what reproduces the equipment that makes that component?

Marx's analysis of social reproduction already distinguished the circulation of commodities from the reproduction of the means of production and of the social relations through which these are organized [124]. In a different analytical tradition, Sraffa's treatment of commodities produced by means of commodities makes interdependence among industries a central feature of production rather than an external complication [172]. Leontief's input–output analysis similarly represents the requirements of one productive sector for the outputs of others [112]. These approaches differ in their accounts of value, distribution, and explanation; none should be described as a derivation of the present generative-relational (GR) formalism. Their shared recognition that productive means are themselves produced supplies a well-established foundation for the problem investigated here.

Contemporary asset management adds an operationally useful perspective. ISO 55000:2024 presents life-cycle asset management and decision-making as activities related to organizational objectives, while reliability-centred maintenance standards specify systematic evaluation of maintenance processes [103, 38]. Such practices provide essential baselines. They can reveal whether the local factory is maintained, but they do not automatically identify every upstream requirement for manufacturing replacement equipment. The question is therefore one of *nested productive reproduction*, not the assertion that conventional industrial engineering has failed to recognize maintenance.

## 8.2 Reproductive Dependencies across Production Levels

Let $B$ be the production boundary and $H = (I_0, \ldots, I_N)$ the observation horizon defined in Section 7. For each required equipment or intermediate-product type $i$, let $\mathcal{Q}_i$ be the set of admissible processes able to supply it at the required quantity and quality. These processes can require other equipment, processed materials, energy services, skilled work, knowledge artifacts, calibration or certification. A *reproductive dependency* exists when maintaining the availability of type $i$ over $H$ requires completion of another qualified productive activity or continued access to an enabling capacity. Importantly, such requirements may be conjunctive: electricity, a machine tool, a controlled environment, and a qualified feedstock may all be necessary for one fabrication process. They may also admit documented alternatives.

Starting with a specified set of activities $\mathcal{T}$, define a sequence of *requirements discovered by expansion*. Set $D_B^{(0)}(\mathcal{T}, H)$ equal to the typed inputs and capacities needed directly for $\mathcal{T}$. For $k \geq 0$, let $D_B^{(k+1)}$ include $D_B^{(k)}$ and the prerequisites of the selected qualified processes that maintain or replace those requirements over $H$. This sequence is indexed by the declared process representation and the admissible substitution policy; it does not assert a unique natural hierarchy of generative levels. An upstream mineral refinery, a testing laboratory, and a contractual service may be reached through different paths, and an alternative source can change the expansion.

A distinction between *dependency discovered* and *dependency satisfied* is indispensable. The fact

DRAFT

Table 10: Nested infrastructure-reproduction requirements. The evidence needed at each level differs from evidence of direct task completion.

| Requirement | Example of a reproductive activity | Necessary evidence beyond local operation |
|---|---|---|
| Consumables | Replenishment of feedstocks, filters or process gases | Qualified source, delivery, material balance and expected future demand |
| Replacement parts | Fabrication and testing of a worn actuator or circuit board | Tool availability, compatibility, yields and supplier capacity |
| Production equipment | Repair or replacement of fabrication and inspection instruments | Capability to maintain tools, metrology and precision components |
| Enabling services | Provision of power quality, cooling, networks or computation | Capacity and reliability of the upstream service infrastructure |
| Knowledge and labour | Updating recipes, software, training and maintenance skills | Tested validity, permitted access and ongoing human and institutional support |
| Ecological conditions | Continued access to water, suitable sites and ecological processes | Independently established ecological constraints and cumulative exposure |

that a database contains a path to a possible supplier does not show that the supplier has spare capacity, meets the requested specification, can deliver during $H$, or is accessible under the relevant institutional rules. Similarly, a circular path within the recorded network can represent actual material recovery, repeat purchases, an accounting convention, or a merely proposed capability. The typed event and provenance description of Section 6 therefore records which relationships are measured, documented, inferred, or hypothetical.

The categories in Table 10 illustrate how an immediate maintenance achievement can depend on successive enabling systems. The categories are analytic examples, not a claim that each facility requires every listed technology.

### 8.3 A Finite Accounting Model of Upstream Requirements

A useful restricted calculation follows the Leontief tradition. Suppose that $m$ output categories can be consistently normalized to specified physical units, and that $A \in \mathbb{R}_{\geq 0}^{m \times m}$ has coefficients $A_{ij}$ denoting units of input type $i$ required per unit gross output of category $j$ over a declared accounting interval. Its entries have units $[i]/[j]$, and matrix multiplication is dimensionally meaningful after the chosen unit conventions have been declared. Let $d \in \mathbb{R}_{\geq 0}^{m}$ denote a specified final output and replacement-demand vector, with the same output-category units, and $z$ the gross production vector. In the fixed-coefficient, linear, unconstrained-capacity approximation,

$$z = d + Az, \qquad (I - A)z = d. \tag{15}$$

The replacement component of $d$ should be estimated separately from the output sold or delivered; otherwise the model conceals consumption of fixed capital behind apparently sufficient finished-product output. Replacement coefficients estimated from wear or failures also depend on the assessment horizon and utilization regime. If a machine is not produced as a continuously divisible commodity, its periodic renewal can require a discrete-event or integer-production model instead. Equation (15) is a restricted aggregate accounting approximation, not a complete dynamic production law.

For a nonnegative coefficient matrix with spectral radius $\rho(A) < 1$, the well-known Neumann representation yields

$$z = (I - A)^{-1} d = \sum_{k=0}^{\infty} A^k d. \tag{16}$$

The successive terms measure direct and higher-order upstream requirements within the *modeled* productive categories. An input absent from the sectoral classification, or a component sourced outside the

DRAFT

observed boundary but coded as an aggregated service, remains absent from the inference. Environmental input–output and hybrid life-cycle inventory methods already investigate such truncation problems; Lenzen, for example, documents omissions of higher-order upstream stages in conventional process inventories [111]. The contribution of a GR-enriched record would have to consist in more accurately qualifying changing alternatives, institutional access, histories and evidence status than these established tools, not in merely rediscovering the inverse in Equation (16).

For a bound involving a norm, mixed physical units require a declared coordinate normalization. Choose a positive reference quantity $s_i$ with the physical unit of commodity $i$ and let $D = \operatorname{diag}(s_1, \dots, s_m)$. The numerical coordinate representations $\widetilde{A} = D^{-1}AD$, $\widetilde{d} = D^{-1}d$, and $\widetilde{z} = D^{-1}z$ are dimensionless. A norm of $A$ computed without specifying such coordinate scales can change merely because one category is expressed in kilograms rather than tonnes. The following bound is therefore stated in normalized coordinates.

**Proposition 2** (A conditional bound on higher-order modeled requirements)**.** *Let $A \geq 0$ and $d \geq 0$, choose the normalization $D$ defined above, and suppose a submultiplicative induced norm satisfies $\|\widetilde{A}\| < 1$. Define $z^{(K)} = \sum_{k=0}^{K} A^k d$. Then the tail* within the modeled category universe *satisfies*

$$\|D^{-1}(z - z^{(K)})\| \leq \frac{\|\widetilde{A}\|^{K+1}}{1 - \|\widetilde{A}\|} \|D^{-1}d\|. \tag{17}$$

*Proof.* Similarity of the matrices gives $D^{-1}A^k d = \widetilde{A}^k \widetilde{d}$. By the norm assumption the normalized Neumann series converges. Submultiplicativity and the triangle inequality give $\|\sum_{k=K+1}^{\infty} \widetilde{A}^k \widetilde{d}\| \leq \sum_{k=K+1}^{\infty} \|\widetilde{A}\|^k \|\widetilde{d}\|$, and summing the geometric series yields Equation (17). □

This proposition is an elementary consequence of established matrix analysis, stated here to clarify the assumptions behind a finite truncation. It is *not* a novel economic theorem. The sufficient condition $\|\widetilde{A}\| < 1$ and the numerical size of the bound depend on the declared normalization and norm; the convergence condition $\rho(A) = \rho(\widetilde{A}) < 1$ is invariant under this coordinate rescaling and does not imply the displayed bound for every norm. The computed tail accounts for higher-order paths among *included* categories only: it provides no upper bound at all for omitted sectors, unrecorded material streams, or erroneous coefficients. An empirical investigator would have to estimate coefficients and their uncertainties, identify omitted categories, and test sensitivity to aggregation, prices, technology changes, and substitution before attaching significance to the calculated upstream tail.

### 8.4 An Illustrative Reproduction Calculation

Consider a strictly hypothetical system with two abstract intermediate-product categories. For one accounting interval let

$$A = \begin{pmatrix} 0.10 & 0.20 \\ 0.30 & 0.10 \end{pmatrix}, \qquad d = \begin{pmatrix} 10 \\ 2 \end{pmatrix}. \tag{18}$$

These coefficients are invented to illustrate the method and should not be read as observations of actual semiconductor or machine-tool production. The matrix has spectral radius approximately $0.345$. The direct-demand vector requires ten qualified units of the first category and two of the second. Accounting for intermediate requirements by Equation (15) gives $z = (12.533\ldots, 6.400\ldots)^{\mathsf{T}}$. Counting through order four produces $z^{(4)} = (12.4902, 6.3466)^{\mathsf{T}}$, already close in this chosen example. Neither result proves that the relevant factories possess the required productive capacities.

This distinction becomes apparent if the second category requires a specialist inspection instrument that is unavailable during the assessment horizon. The matrix calculation may report a nonnegative gross-output vector while actual manufacturing remains infeasible. Likewise, an alternative process may reduce input needs but require a certification the facility does not hold. Capacity ceilings, delivery delays, nonconvex batch production, skilled-labour availability, network failures and institutional authorizations therefore belong to the feasibility analysis in Section 6. Their omission is especially serious when interpreting one feasible quantity vector as evidence of independent system reproduction.

DRAFT

For empirical work, the restricted input–output model is best used to identify a candidate upstream requirement and to define a baseline. A richer activity model can then ask whether sufficient capacity and qualified alternatives exist, which material stocks and tool conditions would be depleted, and what external support remains indispensable. Where the standard model suffices for the decision task, no claim of a GR performance advantage is warranted.

### 8.5 Closure, Substitution, and Continuing External Support

*Closure* has several non-equivalent meanings. A facility may produce its own consumables, repair most equipment types, recycle part of its material stream, or possess the capability to manufacture replacements for selected machines. These claims must not be silently transformed into complete physical independence. Even a very broad productive network uses external energetic and material inputs, operates through finite conversion efficiencies, and exchanges heat and other outputs with its environment. The industrial-ecology literature addresses such interconnections by studying coordinated resource circulation and waste reduction, while also recognizing constraints on real recovery processes [45, 44]. Therefore, a productive loop in a flow diagram is evidence of modeled circulation, not a thermodynamically closed system.

For a declared task set $\mathcal{T}$ and horizon $H$, define *bounded reproductive coverage* of an infrastructure type $i$ when a documented, feasible activity network can restore its specified continuation requirement within $H$, subject to an explicitly stated external-support profile. Stronger coverage includes the qualified tools and resources needed to restore the first-level replacement processes. The requirement can be recursively expanded until a research-defined stopping rule is reached. Rather than assign one scalar "self-sufficiency percentage" to mixed quantities of power, material, specialist labour and software rights, the result can be a typed coverage register identifying (i) what has been restored within $B$, (ii) what crosses $\partial B$, (iii) which alternatives have been demonstrated, and (iv) what remains unknown. A percentage is interpretable only for a consistently defined denominator, such as the mass of a specified recovered material or the share of a counted component catalogue, with its exclusions reported.

Substitution makes the dependency structure conditional on technical and social circumstances. A failed supplier can sometimes be replaced through another qualified supplier, an alternative design, refurbishment, stock reserves, or demand reduction. Another dependency may be technically unique over the relevant horizon. The importance of this distinction is illustrated by OECD's 2025 mapping of semiconductor value chains, which documents specialized equipment, complex multi-tier dependencies and concentrated segments; the report does not claim that AI systems cannot function without a particular individual supplier in every configuration [141]. A reproductive assessment should consequently test *effective substitutability under specified constraints*, not assign structural necessity solely from the presence of a single supplier in a recorded graph.

### 8.6 Reproduction across Unequal Timescales

Maintenance, fabrication, geological extraction, infrastructure construction, skill formation and ecological recovery rarely operate over identical assessment horizons. A spare component may be held in inventory and delivered within days, whereas establishing the process equipment that produces it can require years. The corresponding long-run feasible-production trajectory can depend on investments, qualification procedures and training undertaken before the immediate task campaign. This matters when comparing a facility's apparent autonomy during a short demonstration with the continuity of its productive capacities over a much longer horizon.

A dynamic extension should therefore model inventories and asset condition alongside qualified productive capacity. For an equipment category $i$, a bookkeeping transition can be written

$$K_{i,n+1} = K_{i,n} - \delta_{i,n} + r_{i,n} + m_{i,n}, \tag{19}$$

where $K_{i,n}$ represents service-capable units at interval $n$; $\delta_{i,n}$ denotes units lost from qualified service through wear, failure or retirement; $r_{i,n}$ denotes units restored by qualified repair; and $m_{i,n}$ denotes

DRAFT

qualified additions from replacement production. Each term is measured in service-capable units, not currency. When partial degradation rather than binary qualification is relevant, a vector of condition states or remaining useful life is required. Repair, replacement, and loss rates may depend on workload and earlier interventions. Equation (19) is not a mechanistic reliability law, and recording a unit as restored requires evidence that it meets the applicable technical specification.

Power generation and computing infrastructures enter analogous transitions through different physical variables and operating constraints. Knowledge capabilities and human expertise need yet other descriptions: a software backup cannot stand in for its tested functionality, and a training record cannot be equated with retained practical expertise. Such heterogeneous processes may be jointly necessary even when their dynamics cannot be reduced to the same differential equation. Institutional arrangements further affect the pace and feasibility of replacement through access, contractual rights, investment, certification and occupational organization.

## 8.7 The Historical Organization of Infrastructural Reproduction

The recursive perspective helps clarify a historical-materialist problem that might otherwise appear merely technical. Productive forces are sustained by prior and concurrent acts of production, by inherited machinery and scientific knowledge, and by institutional arrangements that organize labour and resource access. A manufacturer can possess advanced automatic equipment while depending upon publicly supported research, suppliers with very different labour conditions, transnational transport, and energy infrastructures governed elsewhere. The means of production thereby embody social and historical relations that need not be visible from the operator's interface.

There is no necessary one-way political implication. Concentrated ownership of a critical machine or software service may create a potential bottleneck, but the actual exercise of control depends on alternative access paths, contracts, capacities, regulation and collective action. Public or cooperative ownership can alter some of these relations while introducing different coordination or financing constraints. The present section is concerned primarily with the *constitution of productive continuation*; subsequent sections investigate dependency relocation, bottleneck control, and value circulation in greater detail. It would be premature to infer exploitation, legitimacy, or optimal institutional form from the existence of a dependency edge or from a matrix coefficient.

The ecological implication is equally conditional. Local repair and recovery may reduce demand for new components and associated extraction, but may also require additional energy, transport, chemicals and materials. A facility can maintain its machinery while imposing cumulative disturbance on external ecological systems. Its capacity to reproduce its own productive conditions is therefore distinct from the compatibility of that reproduction with ecological continuation. The latter requires spatially and temporally appropriate inventories and ecological evidence, using system boundaries broad enough to include displaced burdens. In the language of the present paper, *technological reproductive viability* and *ecological admissibility* are separate predicates until empirical analysis establishes their relationship.

## 8.8 Research Design for Identifying Reproductive Bottlenecks

A tractable comparative study would select one industrial facility or autonomous laboratory with a documented task campaign and construct four nested descriptions: (a) direct production and maintenance logs; (b) a component–supplier register with qualification and lead-time records; (c) an input–output or hybrid life-cycle account of upstream requirements; and (d) a typed relational register including alternative paths, institutional access and evidence provenance. The fourth description would be evaluated against, rather than presumed superior to, the other three. All models would use a common declared functional output, time horizon, and quality threshold wherever comparison permits.

Researchers could test interruptions involving replacement components, energy quality, calibration services and network availability, with human and ecological safeguards. Outcome variables would include qualified productive output, time to restore service, spare-part and repair demand, external-support

DRAFT

channels in physical units, and changes in continuation readiness. A short-horizon test and a longer-horizon renewal scenario should be distinguished. The research must report which supplier and upstream records are inaccessible, which dependency edges are inferred, and how uncertainty propagates from the inventory or input coefficients into projected feasibility. Historical archives, interviews and contractual sources may be necessary to investigate who exercises control over relevant access decisions; such evidence cannot be replaced by observing machine telemetry alone.

The proposed framework yields definite failure conditions. A hypothesis of systematic dependency concentration would be weakened where qualified substitute pathways diversify and interruptions become less consequential under comparable workloads. A hypothesis of a large reproductive autonomy gap would be contradicted in cases where robust maintenance and replacement consistently restore continuation readiness under stated boundary and support rules. And if the enriched GR registry provides no independently measured benefit over competent supply-chain, reliability, provenance and industrial-ecological approaches, its additional informational cost may not be justified. These tests should be predeclared to reduce retrospective selection of favorable examples.

### 8.9 Implications for the Transformation of Productive Relations

Recursive reproduction reveals a second dimension of productive autonomy: a system may execute tasks and replenish immediate inputs while remaining dependent upon other systems for the reproduction of its own means of replenishment. The input–output calculation gives a restricted quantitative language for higher-order requirements; typed activity records add information about actual qualifications, alternatives, evidence and governance; and asset transitions extend the inquiry beyond a single output campaign. Together they supply a way to study whether autonomous production reorganizes dependencies, and which parts of its reproduction remain socially and ecologically distributed.

The next section examines changes in the *location and organization* of dependencies when autonomous technologies are introduced. In particular, it asks how matched productive outputs can be associated with altered upstream requirements, supplier relations and geographically displaced burdens. The present section establishes the reproduction requirements that such a comparison must trace.

## 9 The Relocation of Dependencies Through Automation

The preceding sections established that operational autonomy depends on the boundary used to identify external support and that the reproduction of productive equipment involves requirements extending beyond the immediate production cycle. This section investigates the historical and material changes in those requirements when automation is introduced. Its object is *dependency relocation*: a change in the composition, location, timing, institutional mediation, or substitutability of the inputs and capacities required to sustain a defined productive activity. The purpose is to distinguish demonstrable reductions in particular requirements from their displacement, concentration, or creation elsewhere in a production network. A relational comparison is developed using typed dependencies, matched functional output, and explicit accounting boundaries. It is intended to complement established input–output, industrial-ecological, and supply-chain methods; it does not assume that automation necessarily increases dependence or ecological burdens.

### 9.1 The Historical Reorganization of Productive Requirements

Industrial machinery transforms both the performance of work and the circumstances under which the work can be continued. The replacement of a labour-intensive process with programmable machinery may reduce the number of person-hours needed on the factory floor, while increasing requirements for precision equipment, specialized materials, power quality, software support, calibration, and maintenance. Some inputs become interchangeable through standardization; others become tied to a small number of qualified processes or owners. Historical materialism directs attention to both the development

DRAFT

of productive forces and the changing relations through which these forces are assembled, maintained, and governed. Machinery embodies accumulated productive capabilities, yet its functioning remains an activity within a wider division of labour and material reproduction [120, 125].

A historically significant change therefore cannot be inferred from a reduction in the locally visible quantity of labour. The correct object is a *productive configuration*: an arrangement of human activities, machinery, intermediate inputs, information, energy, repair capacity, and rights of access sufficient to realize an output under specified conditions. Automation changes this configuration by altering tasks, process qualifications, and the position of dependent activities within wider production relations. A component formerly fabricated inside an establishment may be purchased from a specialized supplier, while a formerly outsourced component may become economical to fabricate on site. The first change externalizes a requirement relative to the establishment; the second internalizes it. Either can improve productivity, degrade resilience, or leave system-wide requirements effectively unchanged, depending on substitutions and counterfactual conditions.

Evidence concerning robot adoption illustrates the importance of level and denominator. Graetz and Michaels' analysis of 17 countries during 1993–2007 reports gains in labour productivity associated with industrial robotics without a statistically significant reduction in aggregate hours worked in its specification; the estimated composition of work nevertheless changed [50]. Acemoglu and Restrepo study US local labour markets and report adverse employment and wage effects in more robot-exposed commuting zones over a different historical period and identification design [3]. These findings are not intrinsically contradictory: national or industry productivity, local employment, and the allocation of production across establishments are different estimands. Neither study provides a direct measurement of contemporary AI factories' recursive component requirements. The contrast demonstrates why an inquiry into relocated dependencies must declare its spatial, institutional, and functional boundaries before interpreting observed changes.

## 9.2 The Forms of Dependency Relocation

In the typed production representation of Section 6, a dependency records a condition for which a productive activity requires another activity, resource, service, or authorized relation. Relocation occurs when an intervention changes the assignment or characteristics of these requirements while the comparison preserves a common productive service. Six mechanisms are analytically useful. *Functional substitution* replaces one requirement with a technically different input or process, as when sensor calibration replaces a repeated manual inspection. *Organizational relocation* changes which establishment or contractor performs an activity. *Spatial relocation* changes where a burden is incurred or an input is produced. *Temporal relocation* moves effort into equipment construction, model training, preventive maintenance, or end-of-life treatment. *Institutional relocation* changes who controls the right to use or reproduce a needed capability. *Risk relocation* changes exposure to uncertain failures, delays, or correlated interruptions. These mechanisms can occur jointly, and their consequences depend upon whether dependencies become more replaceable or more difficult to satisfy.

Table 11 identifies the evidence needed to distinguish these changes. The classifications are analytical proposals, not mutually exclusive empirical categories. A recorded reduction in direct employment at one site, for example, can coexist with new engineering and maintenance work elsewhere, and it does not identify the net employment effect without a defined population and credible counterfactual.

## 9.3 Comparing Matched Production Services

A claim of relocation requires a reference service. Let $s$ identify a productive service with a declared technical specification, tolerances, expected quality, and an observation horizon $H$. Let $j \in \{0, 1\}$ index a baseline and a candidate automated configuration, each serving the same specification. Denote by $q_j$ the number of qualifying service units delivered during $H$; the unit may be accepted parts, completed experiments, or a properly defined flow of production services. Let $k$ identify a typed input or burden category, such as kilograms of a specified material, kilowatt-hours of electricity, qualified labour-hours,

DRAFT

Table 11: Mechanisms of dependency relocation and the observations required to identify them.

| Mechanism | Possible production change | Minimum comparative evidence |
|---|---|---|
| Functional | Labour-intensive operation replaced by a robot-assisted process | Equivalent output quality; direct and embodied inputs; task and error rates |
| Organizational | Replacement components purchased instead of made in-house | Supplier and internal process records; qualified alternatives; contractual conditions |
| Spatial | Processing shifts among sites or regions | Site of actual fabrication; embodied transport; energy mix and local ecological exposure |
| Temporal | Work shifts to equipment fabrication, updates, and maintenance | Asset lifetime; construction and replacement inventories; schedule and failure distribution |
| Institutional | Access moves to proprietary platforms, licences, or controlled infrastructure | Ownership and access conditions; switching rights; service continuity under refusal |
| Risk | Dependencies share a constrained or failure-prone enabling node | Capacity, lead time, substitution time, correlated outage scenarios and recovery evidence |

or a verified component replacement. Let $g$ index geographic or organizational locations and $\tau$ a temporally defined interval. The *per-service dependency intensity* $f_{j,k,g,\tau}$ has units of category $k$ per qualifying unit of $s$ at location $g$ and interval $\tau$, including appropriate amortized upstream contributions under an explicitly declared accounting rule. Define the corresponding total accounting entry by

$$D_{j,k,g,\tau} = q_j f_{j,k,g,\tau}. \tag{20}$$

Equation (20) is an accounting definition rather than an assumption of linear physical response to arbitrary changes in scale. The coefficient $f_j$ is measured or modeled *at* $q_j$, and may itself change with scale, learning, utilization or equipment renewal. Within a single typed component $k, g, \tau$, an exact two-configuration decomposition is

$$D_{1,k,g,\tau} - D_{0,k,g,\tau} = q_0\big(f_{1,k,g,\tau} - f_{0,k,g,\tau}\big) + (q_1 - q_0) f_{1,k,g,\tau}. \tag{21}$$

The first term is a matched-baseline-output comparison of intensities, conventionally interpretable as a configuration or technique effect only if compositional and other confounders have been addressed. The second term attributes the remaining arithmetic difference to the changed number of qualifying outputs at the candidate intensity; it is not, by itself, a causal estimate of rebound. An alternative decomposition using the baseline intensity is equally exact, and the attribution differs when the intensity and quantity both change. Reporting both or using a symmetric index decomposition can reveal this index-number sensitivity.

The entries must retain their physical types. Kilograms of rare-earth inputs, megawatt-hours of energy, and hours of maintenance cannot be summed into an unqualified dependence score. When a scalar comparison is necessary, weights must be justified by a specified engineering, economic, environmental, or normative objective, with uncertainty and non-substitutable constraints retained separately. This safeguards the distinction between physical feasibility, financial cost, ecological admissibility, and socially legitimate production choices developed in Section 6.

### 9.4 Internal Independence and System-Wide Dependence

The effect of shifting a requirement across an establishment boundary can be stated formally. Fix a productive service $s$, its delivered quantity $q$, horizon $H$, disturbance class, and a common list of typed enabling requirements. Let $B$ be the monitored establishment and $W$ a wider production boundary containing $B$ and its selected upstream suppliers. Define $L_B(c)$ as the recorded direct requirement for a chosen dependency category within $B$ under configuration $c$, and $U_{W\setminus B}(c)$ as the corresponding embodied or directly recorded requirement elsewhere in $W$, measured in compatible units and without double-counting. Let $R_W(c)$ be the system-wide requirement, conditional on the selected truncation, allocation, and life-cycle rules.

DRAFT

**Proposition 3** (Local requirement reduction is not sufficient for system-wide reduction)**.** *Suppose the accounting identity $R_W(c) = L_B(c) + U_{W\setminus B}(c)$ holds for a specified input category and fixed qualified service. Without an additional bound on changes outside $B$, $L_B(c_1) < L_B(c_0)$ does not imply $R_W(c_1) < R_W(c_0)$.*

*Proof.* The system-wide difference satisfies

$$R_W(c_1) - R_W(c_0) = \big[L_B(c_1) - L_B(c_0)\big] + \big[U_{W\setminus B}(c_1) - U_{W\setminus B}(c_0)\big]. \tag{22}$$

The first bracket is negative by hypothesis. The second bracket is unrestricted. For an explicit admissible example in a single physical input unit, let the local requirement fall from $10$ to $4$ while the upstream requirement rises from $2$ to $9$. The corresponding system-wide requirements are $12$ and $13$, respectively. Therefore a strictly lower local requirement can coexist with a strictly higher system-wide requirement. The same identity also permits a lower system-wide requirement when the upstream increase is sufficiently small or negative. □

This proposition is an elementary accounting counterexample, not an original theorem about the consequences of automation. It establishes the need for an adequate boundary; empirical determination of the magnitudes and uncertainty of the terms remains essential. When the use of automation changes the composition of inputs, a physically typed vector identity should replace the scalar one. It may then be possible for the candidate to reduce some requirements and increase others without yielding a single dominance relation.

## 9.5 Concentration, Alternatives, and the Structure of Supply

A production system may reduce the number of direct transactions while increasing dependence upon a supplier with few qualified substitutes. Conversely, standardized robotic equipment or open process interfaces may create additional feasible suppliers and reduce effective vulnerability. Supplier count alone is therefore an incomplete measure of dependence. For a necessary component $i$, relevant variables include the number of suppliers that can meet technical qualifications, independent capacity under a specified demand surge, substitution and recertification time, spare inventories, common upstream inputs, transport routes, and the legal conditions for continuing access. Two nominally distinct suppliers may share a crucial sub-supplier, semiconductor design tool, or electricity constraint. An analysis that regards such suppliers as independent would underestimate common-cause exposure.

The OECD's 2025 mapping of semiconductor production describes globally distributed fabrication and related input chains, with concentrated critical segments and growing trade dependencies; its value-chain mapping is evidence of industrial structure, not proof that any particular AI system is autonomously self-reproducing [141]. The IEA's *Global Critical Minerals Outlook 2025* reports that the three largest refining countries collectively accounted for an average $86\%$ of the market in its selected key energy minerals in 2024, up from around $82\%$ in 2020 [65]. The figure concerns selected mineral refining, not a direct measure of the share of the inputs used by AI robots or factories. These sources help motivate the possibility of concentrated enabling inputs, which must be traced through the specific bill of materials, qualified processes, and maintenance dependencies of any examined facility.

The choice to manufacture a component domestically may change exposure to international trade, but does not automatically eliminate concentrated dependencies. In the OECD's *Supply Chain Resilience Review*, broad relocalisation scenarios were projected to reduce world trade by more than $18\%$ and global real GDP by more than $5\%$ while failing to produce consistent gains in modeled resilience [142]. These are scenario-model outputs, not observed historical losses. Their relevance is methodological: the capacity to bring final assembly within a national border differs from the capacity to reproduce all of the enabling equipment and inputs under disruption. The World Bank's analysis of global value chains also emphasizes the historically contingent opportunities and distributional consequences of cross-border specialization [193].

DRAFT

## 9.6 Relocated Ecological Burdens and Material Metabolism

Dependency relocation also changes the geography of ecological intervention. Manufacturing a robotic component may demand water, electricity, extraction, and waste-treatment capacities in upstream regions that are absent from the immediate operational accounts of its user. Environmental-footprint analysis already addresses such problems through multi-region input–output and life-cycle methods. Wiedmann and Lenzen identify the geographical separation of consumption and substantial environmental and social impacts embodied in international trade [190]; environmentally extended input–output approaches such as EXIOBASE provide one class of quantitative baseline for cross-regional accounting [173]. These established instruments are necessary comparators for any claimed advantage of a GR representation.

A robust analysis should link production-system alternatives to separately specified ecological indicators at their appropriate spatial and temporal scales. Electricity consumption measured at the factory meter should be differentiated from upstream generation burdens; total water withdrawal from context-dependent ecological water stress; carbon emissions from local habitat disturbance; and material recovery rates from evidence of ecological regeneration. The uncertainty introduced by missing suppliers, data aggregation, facility utilization, and end-of-life treatment should be made explicit. Lower local emissions can accompany higher off-site impacts, but the opposite is also possible when clean upstream energy or more efficient centralized fabrication replaces a burdensome process. It follows that environmental consequences should be estimated in a comparative life-cycle design, rather than deduced from automation as a category.

The historically materialist question concerns how productive arrangements organize the displacement or reduction of the material costs of production. Whether a factory's production relation is investor-owned, public, cooperative, or commons-oriented can influence supplier choice, disclosure incentives, investment horizons, and whose ecological exposure enters decisions. These institutional effects should be tested rather than treated as fixed consequences of ownership labels. The GR proposal may contribute a trace of how access, authority, and ecological observability evolve together, but its claimed additional insight would require demonstrating improved counterfactual inference or accountability relative to a competent life-cycle and supply-chain baseline.

## 9.7 Identification from Time-Varying Production Histories

Measuring dependency relocation demands more than a before–after comparison of purchase orders. Adoption may coincide with changes in output composition, commodity prices, energy systems, supplier availability, trade policy, demand, and macroeconomic conditions. A facility-level panel should therefore record at least the production specification, realized qualifying output, equipment and software changes, labour by function and contract, intermediate input quantities, maintenance events, transport, supplier qualification, process yields, energy by type, and information or service access. Each observation needs a time and boundary identifier, with an explicit uncertainty or missing-data flag. The dependence on external providers should be recorded separately from their geographical location: an external service can be local, while an internally owned process can rely on distant inputs.

A comparative research design would match product specifications and observation horizons across adopting and non-adopting facilities or construct a transparent within-facility counterfactual. Where appropriate, difference-in-differences or event-study estimators could examine changes in per-unit supplier dependence, repair delays, imported intermediates, embodied environmental burdens, and the dispersion of qualified alternatives. Such designs require assessment of parallel trends, anticipated adoption, changing treatment intensity, spillovers across suppliers, and selection into automation. If similar firms reorganize purchasing for reasons unrelated to AI, the resulting difference is not evidence of an AI-specific mechanism. Interdependence also makes interference across units especially plausible: the automation of a large buyer may change prices or capacities available to comparison firms.

A meaningful evaluation should include at least one setting in which automation demonstrably diversifies supply or reduces full-system material intensity. Such negative cases would prevent the analysis from confusing the concept of relocation with the empirical proposition of increased fragility. The ob-

DRAFT

Table 12: Selected existing evidence relevant to dependency relocation, with distinct units and inference boundaries.

| Source and population | Reported finding | Use and limit for this paper |
|---|---|---|
| Graetz and Michaels (2018); 17 countries, 1993–2007 | Estimated robot-related labour-productivity gain and occupational-composition changes | Historical effects of industrial robots; no direct observation of AI-era material-reproduction networks |
| OECD (2025); semiconductor value-chain mapping | Specialized, concentrated inputs and increasing cross-border interdependence | Identifies candidate upstream dependencies; not a factory-specific feasibility test |
| IEA (2025); selected minerals, 2024 | Top-three-country average refining share 86% | Geographic concentration for selected materials; not an AI-system input fraction |
| OECD (2025); relocalisation scenarios | Modeled world trade reduction above 18% and GDP reduction above 5% | Conditional model projections; not realized restructuring outcomes |
| Wiedmann and Lenzen (2018); international trade | Documented spatial separation of final consumption and embodied impacts | Supports footprint accounting; not proof of marginal automation effects |

jective is to estimate both the direct improvement from automation and the indirect transformations it enables, including possible productivity gains, lower ecological loads, alternative sourcing, and altered bargaining positions.

### 9.8 Relational Interpretation and Research Implications

Dependency relocation connects the reproduction analysis of Section 8 to the coming inquiry into infrastructural control. Automation may alter which productive conditions become critical and which actors can determine access to them. This is a historically situated transformation: a material input becomes a bottleneck only relative to technical requirements, available alternatives, capacities, institutions, and a specified horizon. The present section does not assign a universal direction to that transformation. It instead supplies matched-service accounting, a typology of relocation, an elementary boundary counterexample, and a design for testing concentration and ecological displacement.

Four falsifiable hypotheses follow. First, some autonomous upgrades may increase the proportion of productive requirements mediated by external specialized providers even while lowering local operating support. Second, process modularity and qualification of alternative suppliers may offset or reverse such dependence. Third, facility-level environmental gains may differ from lifecycle gains when major burdens move across space or time. Fourth, access restrictions may influence feasible production independently of the nominal physical availability of equipment or material. Each proposition admits empirical counterexamples and should be tested under declared boundaries and alternative explanations.

The next section examines *control over bottlenecks in productive networks*. Its task is to investigate how dependency positions acquire institutional and economic significance through scarcity, substitutability, ownership, access conditions, and the capacity to interrupt productive continuation. That is a distinct question from whether the dependencies were relocated in the first instance.

## 10 Control over Bottlenecks in Productive Networks

The preceding section established that automation can redistribute productive requirements among establishments, regions, technical stages, and institutions. This section investigates the conditions under which one of those requirements becomes a consequential constraint on productive continuation and the further conditions under which an actor can exercise power through its control. The objective is to distinguish a *technical bottleneck*, defined relative to a productive task and feasible alternatives, from *institutionally exercisable control* over access to the requisite capacity. Historical materialism motivates the inquiry into ownership, appropriation, bargaining, and the reproduction of production relations; production-network analysis supplies well-established accounts of interdependence and shock propagation. The proposed generative-relational (GR) contribution is a time- and boundary-specific representation of the changing

DRAFT

feasible alternatives open to heterogeneous participants. This contribution remains a research proposal until it establishes additional explanatory or predictive value relative to competent network, industrial-organization, and supply-chain analyses.

## 10.1 Bottlenecks within Historically Organized Production

The development of machinery changes which productive requirements are difficult to substitute and which actors must cooperate for production to continue. In Marx's analysis, machinery embodies accumulated labour and scientific knowledge but operates within historically specific relations of ownership, direction, and labour [120, 125]. A technologically crucial piece of equipment can therefore be simultaneously a productive achievement, a costly social investment, and a condition of dependence for producers who lack access to it. None of these descriptions alone establishes the degree of control available to a particular owner or the legitimacy of its exercise. The analysis must distinguish the physical role of an input from the social relations specifying who can use it, under which conditions, and with what possibilities of substitution or collective response.

The classical insight that organizations depend on external resources has been developed systematically by Pfeffer and Salancik [153]. Research on global value-chain governance similarly shows that the complexity and codifiability of transactions and the capabilities of suppliers condition forms of coordination and asymmetry [47]. Production-network economics examines how upstream disturbances propagate and how network structure can shape aggregate outcomes [21, 22]. These approaches already address much of what might loosely be described as infrastructural power. The question for this paper is narrower: whether tracing the *historical evolution of feasible productive continuations*, together with the provenance and contestability of access restrictions, reveals changes that ordinary input–output coefficients or static concentration ratios do not adequately describe.

A bottleneck is relational rather than an immutable property of an object. High-quality metrology may be indispensable to one precision-manufacturing process, yet readily replaced for a product with wider tolerances. The same component supplier may be essential within a two-week repair horizon and dispensable over a six-month transition if alternative designs can be qualified. Demand and product specification also matter: spare capacity at a qualified supplier can convert a single-source dependence into a substitutable requirement, whereas a sudden increase in demand can make formerly adequate alternatives insufficient. It is therefore misleading to infer either dependence or power from the observed rarity of an input without specifying a productive function, capacity, qualified alternatives, and a time horizon.

## 10.2 Identifying Bottlenecks through Feasible Substitution

For consistency with Sections 6 and 9, consider a productive boundary $B$, evaluation horizon $H$, and nonnegative vector $q$ of required qualifying output by product or service type. The vector includes quality and timing requirements in the definition of each type; incomparable services are not added. Let $\theta$ collect capacities, qualified technologies, physical stocks, inventories, failure states, lead times, and exogenous disturbances. Let $a$ specify institutionally permitted access to productive resources and services. A feasible production plan $x$ belongs to

$$\mathcal{F}_{B,H}(\theta, a) = \mathcal{F}^{\text{phys}}_{B,H}(\theta) \cap \mathcal{F}^{\text{access}}_{B,H}(a) \cap \mathcal{F}^{\text{admiss}}_{B,H}, \tag{23}$$

where the intersection is meaningful only when all three sets describe plans in the same typed decision space. The third term represents explicitly chosen safety, ecological, and other admissibility constraints; its normative authority is a separate question. This representation does not imply that all physical or institutional constraints are known. A complete implementation should distinguish observed conditions, estimated conditions, and legally or organizationally contested conditions.

Assume that at least one required service component has $q_k > 0$, so that the minimum below is taken over a nonempty index set. For $q_k > 0$, let $y_k(x)$ be qualifying service delivered by plan $x$ in the same

physical unit as $q_k$. A simple bounded service score is

$$D_q(x) = \min\left\{1, \min_{k:q_k>0} \frac{y_k(x)}{q_k}\right\}, \qquad \Phi_{B,H}(\theta, a; q) = \sup_{x \in \mathcal{F}_{B,H}(\theta,a)} D_q(x). \tag{24}$$

We set $\Phi = 0$ for an empty feasible set, and assume nonnegative qualifying outputs. Thus $\Phi \in [0, 1]$ is a normalized *completion fraction for a fixed output specification*, not a universal index of social value. It is deliberately conservative: under this bottleneck-oriented score, one missing indispensable product can limit completion even when large quantities of other outputs are available. Alternative aggregations may be appropriate for divisible or substitutable outputs, but require a declared service-substitution rule and cannot be inferred from a price vector alone.

The comparison of a capacity loss with the baseline must hold constant the service requirement, policy boundary, horizon, and admissibility rules. If $\theta^{\neg j}$ represents removal or failure of capacity $j$, its *technical criticality* can be described by

$$C_j^{\text{tech}}(B, H, q) = \Phi_{B,H}(\theta, a; q) - \Phi_{B,H}(\theta^{\neg j}, a; q). \tag{25}$$

The expression is meaningful as a nonnegative difference when removing $j$ only reduces available capacity and therefore nests the post-removal feasible set inside the baseline. Replacement, redesign, reserve inventories, and alternative logistics must be admitted in both comparisons on the same terms. The magnitude is a scenario-specific loss of feasible service, rather than a firm's market share or a measurement of its ability to compel another actor's behaviour. A node with many trading links may have small $C_j^{\text{tech}}$ because qualified substitutes exist; a specialized but low-degree process can have high criticality. This distinction limits inferences from graph centrality and concentration alone.

### 10.3 Access Rights and the Possibility of Exercising Control

Physical dependence becomes a potential site of institutional power when some actor can meaningfully change another actor's access to a consequential resource or productive service. The control relation may involve ownership, licensing, procurement specifications, certification, logistics slots, data access, compute allocation, or authority to approve operations. These powers are heterogeneous: the authority to schedule capacity differs from the right to exclude a competitor, and a contractual exclusivity term differs from the capacity to enforce it. A bottleneck may be publicly owned, privately owned, jointly governed, or subject to enforceable public-service duties. Conversely, a large supplier need not control a bottleneck if customers can readily substitute a qualified alternative.

Let $a^{\neg j,r}$ denote a counterfactual in which actor $r$ exercises an *actually available and specified* access lever affecting resource or relation $j$, without simultaneously changing the other model conditions. Define the corresponding *access leverage* as

$$L_{r,j}(B, H, q) = \Phi_{B,H}(\theta, a; q) - \Phi_{B,H}(\theta, a^{\neg j,r}; q). \tag{26}$$

This is a hypothetical capacity to alter productive feasibility through a particular access decision. It is not an estimate of observed coercion, an allegation of abuse, or a measure of normative legitimacy. A credible empirical estimate requires evidence that actor $r$ possesses the relevant decision right, can implement it within $H$, and faces a feasible set of contractual, regulatory, organizational, and retaliatory responses. Costs to the actor exercising the lever, legal remedies, reputation, long-term contracts, and alternative supply arrangements may make an apparent threat unattractive or unenforceable. In practice, access leverage may therefore differ from realized bargaining power and from the economic surplus that an actor can appropriate.

**Proposition 4** (Monotonicity under access relaxation)**.** *Fix $B, H, q$, the physical state $\theta$, and the admissibility rules. If one access regime $a'$ relaxes constraints without restricting any plan formerly feasible under $a$, so that $\mathcal{F}_{B,H}(\theta, a) \subseteq \mathcal{F}_{B,H}(\theta, a')$, then*

$$0 \le \Phi_{B,H}(\theta, a; q) \le \Phi_{B,H}(\theta, a'; q) \le 1. \tag{27}$$

DRAFT

*Proof.* Every feasible plan in the regime $a$ remains a feasible candidate in the regime $a'$. Taking the supremum of the same bounded function $D_q$ over a superset cannot yield a smaller value. The outer bounds follow from the definition of $D_q$ and the empty-set convention. □

This elementary set-inclusion result specifies an accounting discipline, not a novel economic theorem. It does not assert that access liberalization inevitably increases realized production: prices, investment incentives, quality assurance, congestion, safety practices, and changes in capacity may also respond. In such cases $\theta$ or the feasible-plan rules change, and the premise of nested sets must be assessed again. Likewise, removal of a restrictive licence may expand a feasible set without making a single additional unit of output achievable if an independent physical constraint already binds. A structurally prominent actor can possess formal exclusion rights and nonetheless have $L_{r,j} = 0$ for the specified task.

## 10.4 A Constructed Example of Capacity, Control, and Time

Consider a wholly hypothetical facility required to deliver $100$ conforming units of a precision component within a two-week interval. Two qualified external suppliers, $A$ and $B$, can provide at most $80$ and $40$ units in that interval, respectively. Assume additive compatible capacities, no pre-existing stocks, no downstream capacity restriction, and no other admissibility constraint binding. These are intentionally simplifying assumptions, not evidence about a named company. The baseline maximum is $\min\{100, 80+40\} = 100$ units. If a physically indispensable process at $A$ fails, supplier $B$ can deliver only $40$ units, and the short-horizon technical loss is $60$ units. If $B$ fails while $A$ remains available, the loss is $20$ units. These absolute service losses correspond to normalized technical-criticality scores of $0.60$ and $0.20$, respectively, under Equation (25). Both suppliers are visible in the procurement graph, yet their criticalities differ substantially under the same output specification.

Suppose now that $A$ can refuse the order within this interval under a valid contractual option and that no enforceable countervailing rights, substitute inventories, or emergency sourcing are available. The illustrative access leverage of its withholding decision is $60/100 = 0.60$ under the normalized score. If instead the buyer has a binding supply contract enforceable within the interval, the feasible access lever may not exist in the hypothesized form. Even where it exists, the loss of the customer relationship and enforceable penalties may change the supplier's actual incentives. The physical role of $A$ remains unchanged in this comparison, while the institutionally exercisable relation changes.

An alternative supplier $C$ has a potential capacity of $60$ units but requires four weeks for qualification and onboarding. Within two weeks $C$ is inadmissible, so removing $A$ still reduces attainable service to $40$ units. For a subsequent six-week horizon, suppose qualification is complete and $C$ can reliably provide $60$ units during the specified production window without reducing other relevant capacities. Then, after loss of $A$, $B + C$ can meet the matched requirement, and the modeled leverage disappears. This is *horizon dependence*, not proof that the original dependency never mattered. The longer-horizon result is conditional on investment, technical equivalence, available capacity, cost, and qualification actually being achievable.

The example also shows why technical necessity, ownership, organizational rights, and effective power require separate measurements. The possibility that the controlling party can profit from denial or secure favourable contractual terms would require a bargaining model with outside options, surplus, incentives, and enforcement; it cannot be calculated by equating $L_{r,j}$ with a markup or wage effect. Furthermore, expanding alternate supply may redistribute environmental burdens and labour requirements. Technical substitution therefore remains only one part of a historically materialist evaluation.

## 10.5 Observed Concentrations and Their Evidential Limits

The contemporary semiconductor value chain offers a relevant empirical setting without establishing that a particular actor has exercised coercive control. The OECD's June 2025 mapping study describes globally distributed and interdependent production segments, with concentrated critical inputs and specialized regional capacities [141]. Its December 2025 study of wafer-fabrication capacity reports that

DRAFT

Table 13: Constructed capacity example: matched demand of 100 conforming units per assessment window. Results are conditional and do not represent empirical supply-chain observations.

| **Scenario** | **Maximum qualifying output** | **Interpretation** |
|---|---|---|
| Two-week baseline; $A = 80$, $B = 40$ | 100 units | Required service can be completed, with 20 units of nominal unused capacity |
| Two-week loss of $A$ | 40 units | Loss is 60 units; normalized technical criticality $C_A^{\text{tech}} = 0.60$ |
| Two-week loss of $B$ | 80 units | Loss is 20 units; normalized technical criticality $C_B^{\text{tech}} = 0.20$ |
| Six-week qualified replacement; $B = 40$, $C = 60$ | 100 units | Loss of $A$ need not restrict service if replacement conditions are actually met |

Table 14: Contemporary evidence relevant to potential bottlenecks, with the inference that remains unestablished.

| **Evidence and reference period** | **Supported observation** | **Additional evidence needed for a power claim** |
|---|---|---|
| OECD (December 2025), worldwide wafer capacity | Five economies: 87% of in-production capacity in September 2025; ten firms: about half | Process-node compatibility, replacement lead times, access rights, capacity available to a specified buyer |
| IEA (2025), selected mineral refining in 2024 | Mean top-three-country share of 86% | Material-specific qualified alternatives, inventories, contracting, ownership and operational impact |
| FTC staff study (January 2025), selected AI partnerships | Documents categories of access, exclusivity and switching-cost arrangements; notes potential competition concerns | Enforceable levers, use of those levers, competitive effects and legally warranted findings |
| OECD supply-chain scenarios (2025) | Broad relocalisation scenarios show modelled output and trade costs without consistent resilience gains | Case-specific causal effects of particular sourcing and diversification policies |

China, Chinese Taipei, Korea, Japan, and the United States together accounted for approximately 87% of global *in-production* wafer-fabrication capacity as of September 2025, while the ten largest manufacturing companies accounted for around half of worldwide capacity [143]. These are estimates for wafer production capacity at a particular time. They are not concentration indices for all semiconductor design, lithography, packaging, or assembly processes, and capacity at one process node is not necessarily substitutable for another. Hence a geographic or corporate share cannot by itself establish the loss of feasible AI-factory production after an identified disruption.

The IEA's *Global Critical Minerals Outlook 2025* reports that the top three refining countries averaged approximately 86% of production across selected key energy minerals in 2024, compared with about 82% in 2020 [65]. Concentration can indicate exposure to supply shocks and the importance of diversification, yet its relevance to a particular industrial robot depends on the quantities, material grades, inventories, and replaceability of the inputs used in that robot's components. Treating the 86% figure as the probability of a supply interruption or the proportion of inputs controlled by one actor would be a category error.

Institutional evidence reveals a different kind of bottleneck. In January 2025, the US Federal Trade Commission published a staff study of investments and partnerships among three large cloud service providers and two prominent generative-AI developers. The report describes relevant consultation, control, exclusivity, computing-resource, and switching-cost arrangements and identifies possible competition concerns [39]. It is evidence about particular reported arrangements and *potential* effects, not an adjudication that each partnership unlawfully restricted a market or a quantitative measure of the productive leverage of any named firm. It nevertheless illustrates the need to distinguish contractual access to computational inputs from physical chip supply and from demonstrated adverse behaviour.

DRAFT

The OECD *Supply Chain Resilience Review* likewise cautions against equating domestic production with resilience. Its 2025 modelling suggests that broad relocalisation could reduce world trade by more than 18% and world real GDP by over 5%, without consistently improving resilience across the simulated economies [142]. These are scenario outcomes under model assumptions, not observed costs of a completed relocation. They warn that an intervention intended to eliminate one geographically visible bottleneck may create other constraints or forego gains from specialized exchange.

## 10.6 Changing Bottlenecks through Investment and Institutional Reorganization

A productive bottleneck can change even without a physical failure. New standards may make components interoperable; process innovation can lower dependence on a particular material; qualified stockpiles can bridge temporary disruption; and an agreement to share testing equipment can broaden otherwise restricted access. Conversely, a change in licensing terms, export controls, procurement eligibility, or inspection requirements can narrow the set of lawful or feasible alternatives. Some restrictions protect quality, safety, environmental interests, or rights of affected communities. A reduction in the feasible production set is therefore not automatically an injustice; the institutional grounds and distribution of consequences must be examined separately.

Reconfiguration often involves historical investments whose benefits depend on the responses of other actors. A new fabrication facility requires not only machinery but specialized workers, input chemicals, metrology, power infrastructure, credible demand, and opportunities for component qualification. Investments that reduce the number of physical failure points can also impose large fixed costs or reinforce dependence on a narrower ecosystem of tool maintenance and design interfaces. The relevant outcome is the changed structure of feasible productive continuation, not simply the increase or decrease in the number of suppliers. Existing production-network studies provide a baseline for analyzing the propagation of shocks and endogenous adjustments [22, 9].

A full GR representation could record changes in who created, validated, financed, and maintained alternative production pathways. Such histories may help distinguish a resilient network from one that appears diversified on paper but relies on the same upstream equipment, shared single-point service, or common authorization. The advantage, if any, should be demonstrated through an observable improvement in identifying critical dependencies, estimating feasible alternatives, reconstructing decisions, or evaluating appropriate remedies. Recording longer trajectories without increasing the accuracy or contestability of analysis would add documentation cost rather than a demonstrated scientific benefit.

## 10.7 Empirical Identification of Structural Control

A credible case study must specify the productive output at risk, the observation period, the actors able to grant or refuse access, and the plausible substitutions. Candidate data include product-level bills of materials, qualification histories, supplier capacity and order books where lawfully accessible, audited contracts, equipment-maintenance logs, inventory coverage, cross-ownership records, shipping lead times, and records of actual interruption or renegotiation. Data access may itself be asymmetric, and supplier identities or commercial terms may be confidential; the research must respect these constraints rather than infer undisclosed powers from sparse public information.

Three distinct empirical questions should be tested. First, is a candidate input a binding technical constraint after inventories and qualified substitutes are considered? Second, does an identifiable actor possess a legally and practically exercisable decision right affecting that constraint? Third, are observable terms, quantities, prices, delays, or opportunities different because that actor actually exercises the right? The first question can be studied through network feasibility and stress tests; the second requires institutional and legal documentation; the third requires a causal design or carefully bounded counterfactual, with attention to the endogenous behaviour of buyers, suppliers, governments, and competitors. A high $L_{r,j}$ in a constructed model answers neither the second nor the third question by itself.

Where data permit, an event study around an unanticipated capacity outage, an access-rule change, or a supplier qualification could compare production interruptions and substitution costs against matched

DRAFT

facilities. Assumptions about common trends, spillovers, anticipation, differences in product specifications, and changes in the composition of output require explicit checks. More importantly, a positive result must be compared with established concentration measures, network criticality, and standard industrial-organization explanations. A useful negative result would show that an apparently concentrated supplier has little feasible leverage because customers have inventories, enforceable rights, or rapidly qualifying alternatives. Another would show that publicly supported shared infrastructure improves access but fails to change completed production because physical capacities remain insufficient.

The evaluation should extend to labour and ecology. Qualified substitution may depend on additional technicians, outsourced maintenance, longer transport distances, higher energy consumption, or resource-intensive fabrication. Conversely, shared access and standardization may permit more efficient utilization and reduce redundant material investment. Differences in wages, employment composition, environmental burdens, and repairability require separate measurement and cannot be inferred from the access leverage statistic. A historical-materialist explanation should address how production relations influence these outcomes without assuming a universal political or ecological effect from a particular ownership form.

### 10.8 Productive Power, Legitimacy, and the Circulation of Value

The analysis supports a carefully bounded interpretation of infrastructural power. A productive condition acquires strategic significance when its interruption removes otherwise feasible production trajectories. A specific actor can potentially exercise leverage when it can affect access to that condition, and the affected parties lack sufficiently rapid or economical alternatives. Actual authority and realised influence additionally depend on law, organization, bargaining, information, and responses by others. Concentration, technical centrality, ownership, power, and legitimacy are therefore analytically separable variables whose relationships have to be investigated.

There is also a normative boundary to the model. Ecological requirements, labour protections, scientific validation duties, and public authorizations are not merely "frictions" to eliminate in the name of production. The socially appropriate feasible set can be narrower than the physically possible set for reasoned and reviewable purposes. An actor's ability to veto an intervention may protect legitimate interests even if the veto lowers $\Phi$. A welfare loss from exclusion cannot be determined merely from foregone production volume, and the rights of nonhuman life cannot automatically be reduced to ownership of a scarce productive input. The model describes conditional continuations; evaluation of their desirability belongs to independently justified ecological, social, and institutional inquiry.

The distinctive hypothesis is consequently modest but testable: as AI-mediated production reorganizes machinery, information, and supply, strategically consequential bottlenecks may change location and institutional form, and a provenance-sensitive account of these changes may improve the explanation of who can continue production, who bears adjustment costs, and who participates in decisions about feasible alternatives. Counterexamples in which bottlenecks dissolve through modularity, public provision, cooperation, or successful substitution are integral to the programme. The next section examines *appropriation within heterogeneous value circulation*, separating the production and distribution of different forms of value from the narrower capacity to interrupt a productive task.

## 11 Appropriation within Heterogeneous Value Circulation

The preceding section distinguished a productive bottleneck from the enforceable possibility of exercising control over it. This section investigates what can be appropriated through historically organized production relations and how the resulting benefits, burdens, and capacities circulate among heterogeneous participants. The inquiry is necessary because the ability to complete an output, the monetary receipts arising from its sale, control over future productive opportunities, the usefulness of scientific knowledge, and the continuing capacities of ecological systems are analytically different matters. A model of productive autonomy that measures only the volume or price of completed outputs can therefore

DRAFT

misdescribe the social and ecological conditions of productive reproduction. The purpose here is to construct a disciplined account of *appropriation* that retains the distinctions among production, entitlement, transfer, capture, investment, and regeneration; its empirical utility must be assessed against established distributional accounting, input–output analysis, political economy, and institutional research.

## 11.1 The Historical Organization of Appropriation

In Marx's account of capitalist production, the labour process produces use-values while the valorization process organizes the production and appropriation of value and surplus value within historically specific relations of property and wage labour [123, 122]. This distinction already resists reducing the significance of production to physical output: useful goods, monetary realization, and the ownership of the resulting product enter the analysis in different ways. The subsequent development of large-scale machinery further modifies the conditions under which social knowledge and accumulated labour function as productive forces [120, 125]. These are established themes in the critique of political economy, and the present paper does not claim to introduce them.

Nor can every transformation of circulation be inferred from a change in technical capacity. An industrial robot can reduce direct labour per manufactured unit without determining whether the associated cost saving accrues to a worker, consumer, supplier, platform, owner, or public authority. The effect depends on wages, bargaining, pricing, property rights, product competition, access to scarce inputs, and the distribution of risk. Conversely, publicly supported scientific research can improve a firm's productive options without conferring legal ownership of the downstream equipment on the research community. The resulting configuration requires evidence about the actual institutions through which a benefit is realized and assigned. Historical materialism contributes an inquiry into how these institutions are reproduced and contested, rather than a presumption that one allocation necessarily follows from a given technology.

An expanded view of production also requires attention to its enabling conditions. Social-reproduction scholarship examines the care, subsistence, education, and institutional arrangements that sustain labouring and social capacities; ecological political economy examines material exchanges and transformations that support and constrain productive continuation [43, 40]. These perspectives identify productive contributions and dependencies that may be important even when no directly corresponding market transaction occurs. A conceptual distinction between such contributions and exchange value does not settle their ethical or political valuation. It instead specifies what must be investigated separately.

## 11.2 The Distinct Outcomes of Productive Activity

For a bounded set of productive activities $\mathcal{E}_{B,H}$ within system boundary $B$ and evaluation horizon $H$, consider five domains of outcome. First, *physical use outcomes* include conforming manufactured quantities, delivered energy, repaired equipment, and usable services, recorded in their appropriate technical units. Second, *monetary outcomes* include sales receipts, purchased inputs, wages, taxes, financing costs, and monetary claims on residual income. Third, *productive-capacity outcomes* include maintained machines, trained competencies, access to qualified designs, and rights to operate or modify infrastructure. Fourth, *epistemic outcomes* concern the reliability, accessibility, provenance, and interpretability of newly produced knowledge. Fifth, *ecological outcomes* concern changes in material stocks, exposure, habitat processes, and capacities for regeneration or continued ecological activity. These domains can be coupled causally without sharing an intrinsic common unit.

A simple typed outcome record for productive event $e$ is

$$\mathbf{v}_e = \left(\mathbf{q}_e^{\text{phys}},\ \mathbf{m}_e^{\text{money}},\ \mathbf{k}_e^{\text{capacity}},\ \mathbf{s}_e^{\text{epistemic}},\ \mathbf{z}_e^{\text{eco}}\right), \tag{28}$$

where each component has a declared measurement convention and may be partially observed. The tuple represents a family of records, not an element of an ordinary vector space in which its coordinates can

Table 15: Heterogeneous productive outcomes and their evidential boundaries. The domains represent distinguishable inquiries rather than interchangeable quantities.

| **Outcome domain** | **Possible observed indicator** | **Potential claim or control relation** | **Important limitation** |
|---|---|---|---|
| Material production | Conforming units, usable capacity, energy or material balances | Ownership of output, access to equipment, procurement rights | Physical output alone does not determine income or legitimate ownership |
| Monetary circulation | Revenue, wages, payments, operating surplus, taxes | Contractual shares, ownership income, fiscal claims | Price and accounting value added are not measures of all social or ecological contributions |
| Reproductive capacity | Equipment availability, repair skills, reserve components, future access | Maintenance duties, design rights, licences, qualification | A recorded contribution does not prove the capacity can be reproduced independently |
| Scientific knowledge | Validated results, reusable protocols, provenance, accessibility | Authorship, licences, data stewardship, research rights | Citation counts or openness alone do not establish epistemic quality or equitable benefit |
| Ecological continuation | Material stocks, habitat state, exposure and recovery indicators | Stewardship duties, restrictions, participation and remedy | Nonhuman capacities cannot be inferred from market prices or treated as property by default |

be added indiscriminately. Some ecological and epistemic outcomes are qualitative or ordinal; others are quantitative but noncommensurable without defensible conversion rules. A scalar social-value index would require an independently justified normative aggregation procedure, including whose interests count, what substitutions are admissible, and who may contest the weights. Nothing in Eq. (28) supplies such a procedure.

The distinction also matters for the term *value added*. In national accounts, monetary value added is defined through accounting conventions concerning output and intermediate consumption. The OECD's Trade in Value Added (TiVA) programme uses inter-country input–output tables to estimate origins of monetary value added embodied in trade, final demand, and global value chains [139]. Its January 2026-revised 2025 edition reports coverage of 80 economies and 50 economic activities over 1995–2022 [144]. These are valuable empirical instruments for investigating monetary production networks; they do not provide a direct measure of ecological generativity, recognition of contributions, or normative entitlement to benefit. Likewise, use-value in Marxian analysis should not be silently equated with either the accounting category of gross value added or a universal welfare score.

### 11.3 Tracing the Relations of Contribution and Entitlement

The typed dependency model of Section 6 represents which activities rely on inputs, knowledge, labour, and infrastructure. Appropriation requires further relations describing how participants are entitled to access or exclude, receive payments, bear obligations, or retain capacities. Let $\mathcal{R}$ denote identifiable actors or institutional roles, with appropriate safeguards for privacy and commercially sensitive information. Let $\mathcal{K}$ denote categories of claim, including use, modification, distribution, exclusion, oversight, and contestation. A provisional rights record is

$$\mathcal{P}_{B,H} \subseteq \mathcal{R} \times \mathcal{K} \times \mathcal{A} \times \mathcal{J}, \tag{29}$$

where $\mathcal{A}$ is the set of relevant productive resources or activities and $\mathcal{J}$ identifies the legal, contractual, organizational, or customary basis of a putative claim. Such a record cannot establish legal validity through inclusion in a database. Verification requires competent documents, relevant jurisdiction, enforceability, and the responses of affected parties. A claim over scientific data may coexist with conditions restrict-

DRAFT

ing disclosure of personal, sensitive ecological, or traditional knowledge. The absence of a recorded exclusion right also need not imply unrestricted use.

For monetary transactions, define $T_{re}$ as the signed net monetary transfer to actor $r$ attributable to event $e$ under a specified accounting boundary. Consistent treatment of payer and recipient yields the bookkeeping identity

$$\sum_{r \in \mathcal{R}_B} T_{re} + T_e^{\text{outside}} = 0, \tag{30}$$

where $T_e^{\text{outside}}$ records flows crossing the boundary; the identity holds only when every recorded transfer has matching counterpart entries, with the same currency and valuation date. It is not a theory of who *created* value. It is also unsuitable for comparing unconverted currencies, uncompensated labour, nonmarket ecological change, or stock changes as if they were transfer entries. If the analysis instead concerns cash operating surplus, the definition of revenue and allowable expenses must be declared independently; the residual should not be mislabeled Marxian surplus value without an accompanying theoretical derivation.

For knowledge or productive infrastructure, an access relation may matter even where the corresponding monetary transfer is zero. A permissive licence can enable many downstream users to develop compatible devices; a restrictive interface or quality certification can limit who may sell a conforming component. A nonexclusive knowledge resource may permit simultaneous uses, whereas laboratory time, electricity, and physical machinery involve constrained capacities. For this reason the study should measure monetary flow, physical flow, access, and rights using linked but typed records. Provenance tools can help trace the sequence of events through which a claim arose, but provenance is evidence about an asserted relation, not an automatic certification of credit, ownership, justice, or ecological acceptability [196, 191].

### 11.4 Generative Perspectives and the Circulation of Benefits

A generative perspective on economic relations attends to how capacities to create value arise, circulate, and remain available to those participating in their generation. Ron Eglash's account of *generative justice* is an important intellectual inspiration for investigating the circulation of benefits in relation to their productive sources, including ecological, labour, and social processes [31]. The present paper draws upon the value of that perspective while developing its own, provisional GR representation of productive events, dependency changes, institutional claims, and heterogeneous continuations. Neither the GR formalism nor the specific mathematical objects proposed here are attributed to Eglash's theory.

A technologically successful production cycle can leave its enabling capacities in very different conditions. Monetary receipts might finance repair, worker development, ecological monitoring, and wider access to knowledge; alternatively, the associated claims may permit the withdrawal of cash returns while qualified equipment, supply relations, or ecological conditions deteriorate. Both patterns can occur within more than one legal ownership form, depending on governance, financial requirements, resource availability, and performance. Evaluating them requires observing how the benefits of an activity are distributed *and* whether the activity enables future contributions under ecologically and institutionally admissible conditions.

Commons scholarship similarly demonstrates that shared resource governance depends on rules, monitoring, participation, and the characteristics of the resource, rather than a simple opposition between private and public ownership [146]. The question for autonomous AI production is therefore whether particular institutional arrangements support productive continuity, reciprocal access, informed participation, and repair of demonstrable harms. A shared model or publicly funded data repository can be useful without automatically creating collective authority over the energy infrastructure and machinery required to deploy it. A cooperative enterprise can distribute voice while still confronting restrictive licensing, scarce finance, or environmental constraints. A public organization can maintain essential infrastructure while imposing exclusionary access rules. These variations are objects for comparative research.

DRAFT

## 11.5 Why Aggregate Output Cannot Identify Appropriation

The feasible-service score $\Phi$ defined in Section 10 describes completion of a specified productive requirement. It gives no distribution of the corresponding monetary or nonmonetary outcomes. The distinction can be expressed through an elementary identification observation. For a fixed production event, suppose monetary receipts are distributed to $n \geq 2$ recipients in a nonnegative vector $\mathbf{t} = (t_1, \ldots, t_n)$ and observed only through their aggregate $R = \sum_r t_r$. Different distributions can produce the same aggregate, even when the physical output and recorded revenue are unchanged.

**Proposition 5** (Non-identification of distributive incidence from aggregate receipts)**.** *Let $n \geq 2$ and let a nonnegative distribution vector* $\mathbf{t}$ *have two strictly positive coordinates. If the researcher observes only $R = \sum_{r=1}^{n} t_r$, then $R$ does not uniquely identify* $\mathbf{t}$*. In particular, there exist distinct nonnegative* $\mathbf{t}'$ *with $\sum_r t'_r = R$.*

*Proof.* Choose indices $i \neq j$ with $t_i, t_j > 0$ and any $\epsilon$ satisfying $0 < \epsilon < t_j$. Define $t'_i = t_i + \epsilon$, $t'_j = t_j - \epsilon$, and $t'_k = t_k$ otherwise. Then $\mathbf{t}' \neq \mathbf{t}$, all coordinates remain nonnegative, and the sum is unchanged. Thus observing $R$ alone does not identify its allocation. □

The proposition is elementary accounting algebra, not an original theorem about exploitation or ownership. It does not determine which allocations are legally possible, whether output quantities remain unchanged following a real institutional reform, or whether the distributions are normatively justified. It illustrates why a claim about appropriation cannot be established from aggregate industry output, total wages, or global production value without further evidence about transfers, entitlements, and counterfactual production responses. An analogous non-identification problem concerns nonmonetary contributions: an observed product may have depended on several research and maintenance activities, but its completion alone does not identify the causal marginal contribution of any one participant. Joint indispensability and interactions can make a unique allocation unattainable without additional modeling or conventions.

## 11.6 A Hypothetical Comparison of Circulation Arrangements

Consider an explicitly constructed autonomous precision-manufacturing facility. In each of three imagined institutional arrangements, the facility sells $100$ qualifying units at a common price of $2$ accounting units per unit, earning $200$ monetary units in a period. Let independently purchased intermediate inputs cost $80$ and contracted direct compensation cost $60$. Define the remaining illustrative monetary amount of $60$ as a residual *before* any additional taxes, financing charges, depreciation beyond specified costs, or other possible obligations; it is a hypothetical bookkeeping remainder, not an empirical profit rate or a Marxian surplus-value estimate. Suppose each arrangement initially uses identical physical equipment, achieves identical conforming output, and faces identical environmental admissibility requirements. Only the disposition of this illustrative remainder varies.

DRAFT

Table 16: Hypothetical allocation of a 60-unit accounting remainder from identical production. Columns express different monetary destinations, not measured outcomes or judgments of institutional merit.

| Arrangement | Distribution | Repair fund | Worker fund | Ecology fund | What evidence is still needed |
|---|---|---|---|---|---|
| A: investor-return priority | 40 | 15 | 5 | 0 | Actual ownership claims, residual risks, maintenance and ecological effects |
| B: shared-capacity priority | 10 | 25 | 15 | 10 | Cooperative or public rules, oversight, adequacy of allocations |
| C: ecological-reserve priority | 10 | 15 | 5 | 30 | Ecological mechanisms, measured outcomes, monitoring and enforcement |

Each row in Table 16 sums to the same 60 accounting units. The invented allocations differ in who receives or controls a fund, not in observed technical performance. A higher repair allocation might increase future equipment availability, but only if qualified components and repair work can actually be purchased and deployed. A larger ecology fund might pay for effective monitoring or restoration, yet monetary earmarking alone does not show ecological recovery. A worker fund could support capability formation, or fail to improve conditions if its governance and use remain inaccessible to those concerned. This example cannot establish the environmental superiority of a particular ownership arrangement; it clarifies the different causal links that a comparison would need to measure.

Furthermore, the identical-output assumption may become untenable once an allocation policy changes. Investment choices, skill retention, labour supply, product prices, upstream supplier agreements, ecological safeguards, and stakeholder confidence can respond. A valid comparison would estimate these responses, account for financial obligations and opportunity costs, and track at least one complete equipment maintenance or replacement interval. Such a design could yield null, adverse, or mixed findings. The appropriate question is whether an arrangement *actually sustains productive and ecological capacities under its constraints*, rather than whether its accounting categories appear more generous.

## 11.7 Empirical Measurement and Competing Explanations

Three classes of evidence should be brought into the same study without conflating their units. The first concerns monetary allocation: national accounts, the OECD TiVA series, enterprise accounts, wage records, tax data, and reliable contractual information. The January 2026 revision of the TiVA 2025 edition spans 80 economies and 50 industry categories for 1995–2022; its coverage supports cross-economy analyses of accounting value added, while its industry aggregation cannot identify the ultimate distribution of a particular AI-controlled factory's income among individual contributors [139, 144]. Linked enterprise or administrative records would be needed for firm- or role-specific incidence. Where unavailable, the limits on identification should be made explicit rather than replaced with fictional allocations.

The second class concerns scientific and infrastructural contribution: validated experiments, maintained components, code releases, licences, procurement terms, repair logs, data provenance, and actual access to downstream productive opportunities. A cited research paper documents a scientific contribution but does not identify the downstream commercial share attributable to it. Citation counts, patent filings, licences, and model releases are partial and institutionally selected indicators. Their changes can reflect documentation practices or strategic filing behaviour rather than substantive increases in knowledge generation. Existing accounting and innovation studies should be used as baselines for evaluating whether GR trajectory records explain additional variation.

The third class concerns material and ecological conditions: physical material-flow accounting, en-

DRAFT

ergy and water balances, life-cycle assessments, and independently observed ecological indicators. Environmental and social footprint research already addresses indirect effects along international trade chains [190]. Any GR extension needs to demonstrate why explicit histories of changing access and relational dependence improve inference over well-implemented environmental input–output or life-cycle methods. This comparison should include data gaps, uncertainty propagation, appropriate spatial scales, and possible conflicts between short-term output and long-term ecological continuation.

A prospective empirical design could compare production arrangements over matched product specifications, safety requirements, and evaluation horizons, while recording the distribution of cash flows, retained maintenance capacity, workers' substantive participation, scientific reuse, and ecological exposure. The design would need quasi-experimental or matched counterfactual methods where institutional choices are endogenous. Product mix, existing wealth, market concentration, public subsidies, geographic resource endowment, regulation, and capital constraints are potential confounders. Evidence that allocation changes merely follow underlying productivity changes, rather than causing them, would qualify the claimed mechanism. Evidence that more open knowledge circulation increases productive access without improving ecological outcomes would likewise challenge any overly broad hypothesis about beneficial circulation.

Finally, normative assessment remains separate. An ecological constraint can legitimately prevent otherwise profitable production, and a refusal to release sensitive data can protect valid rights even when it reduces measured scientific reuse. A monetary contribution to ecological repair may be insufficient where an intervention threatens irrecoverable habitats. Decisions about rights, consent, environmental thresholds, and remedy require substantive justification and accountable procedures rather than an optimised transfer share. Historical materialism directs attention to the material and institutional conditions behind these decisions; it does not itself convert every heterogeneous outcome into one commensurable value quantity.

### 11.8 Implications for the Reproduction of Productive Relations

The analysis yields a provisional research claim: as AI-mediated autonomy reorganizes productive dependencies, it can also change the relations through which contributions are recognized, productive capacities are retained, benefits circulate, and burdens are assigned. The direction is historically and institutionally contingent. Access to shared knowledge may enlarge the set of feasible future producers, while dependence on specialized equipment can maintain concentrated control; both may occur simultaneously. Different regimes of ownership or stewardship can finance maintenance, capability formation, ecological safeguards, or cash distribution with different consequences. Establishing those consequences requires evidence beyond the nominal ownership form and beyond immediate production efficiency.

The distinct contribution proposed here is an account of how an observed productive event enters multiple, linked but heterogeneous trajectories: physical transformation, economic circulation, reproduction of skills and infrastructure, formation of knowledge, and ecological continuation. This account should be retained only to the extent that explicit provenance and conditional dependence improve explanations or predictions compared with established monetary accounting, global value-chain analysis, commons scholarship, and industrial ecology. The next section examines the reorganization of living labour and social reproduction, where these questions become especially consequential for livelihood, participation, knowledge transmission, and the reproduction of human capacities.

## 12 The Social Reproductive Consequences of Labour Reorganization

The preceding section examined how the results of production may be appropriated, circulated, and retained within heterogeneous productive relations. This section investigates the consequences of AI-mediated labour reorganization for the continuation of human productive and social capacities. The historical-materialist problem concerns the changing division of work, the institutions through which people gain access to means of subsistence, and the often incompletely recorded activities through which

DRAFT

productive capacities are maintained between production cycles and across generations. The fact that a machine can execute a task does not by itself determine who remains employed, whose working time changes, where necessary work is performed, or whether workers and their dependants acquire greater opportunities for education, care, participation, and security. Those consequences require distinct explanations and evidence. The section therefore connects the dependency and appropriation models developed in Sections 10 and 11 to task allocation, working conditions, socially reproductive activities, and institutional comparison.

## 12.1 The Reproduction of Labouring Capacities within Historical Production Relations

Marx's analysis of capitalist reproduction makes the renewal of labour-power a condition for the continuation of the production process; it also distinguishes the purchase and use of labour-power from the reproduction of the social relations in which these processes occur [121, 122]. Such an analysis directs attention to wages and subsistence, the organization of time outside the workplace, and the accumulation of skills and knowledge through which future labour becomes possible. It does not entail that household arrangements, state provision, or the position of workers are constant across historical settings. The social conditions of productive reproduction vary with property relations, labour institutions, migration, schooling, public services, technical organization, and the distribution of unpaid work.

Social-reproduction scholarship makes explicit the dependence of economic production on care, domestic provisioning, education, health, and relations of interdependence that are incompletely represented by market accounts [43, 73]. These activities cannot be treated simply as inventories of unpaid inputs available for automation. They support the lives of people whose needs and interests extend beyond their role as productive labour, and their character changes with demographic, household, and institutional conditions. Different scholars disagree on the conceptual boundaries between social reproduction, care, production, and accumulation; the present study uses the term operationally for processes through which the capabilities, wellbeing, and material conditions of people are continued and renewed. No assumption that a unit of care can be exchanged for a unit of manufactured output is required.

AI-mediated production introduces possible changes in this arrangement at more than one level. Production tasks may be transferred from workers to equipment. Coordination and supervision may be intensified through algorithmic management. Workers may acquire new maintenance, verification, and safety responsibilities. The resulting distribution of income, working hours, training, and access to essential services depends on institutions beyond the technical boundary of the machine. A production system can remain technically operational while the social infrastructure that trains its operators, sustains household life, or enables workers to participate safely is weakened. Conversely, changes in production can expand the resources available for education, care, and leisure when appropriate institutions and material provision exist. Neither direction follows from automation alone.

## 12.2 Task Reallocation and the Indeterminacy of Employment Outcomes

Technological exposure, adoption, task substitution, working-time change, unemployment, and aggregate welfare are different observables. In task-based labour economics, automation can displace labour from specific tasks while productivity effects and the formation of new tasks may sustain or increase labour demand elsewhere [2]. A single occupation typically contains tasks with heterogeneous requirements. The technical feasibility of automating one task does not demonstrate that its associated occupation disappears, nor does the preservation of the occupation establish that work intensity or bargaining power remains unchanged. Moreover, organizations select among technically feasible systems according to costs, finance, production strategy, managerial preferences, safety requirements, and applicable labour institutions.

Published estimates illustrate the dependence of observed outcomes on the technology and population studied. Acemoglu and Restrepo's analysis of industrial robots across US commuting zones found negative effects on local employment and wages for the historical period and design examined [3]. Graetz and

DRAFT

Table 17: Selected empirical observations relevant to labour and social reproduction. Populations, observation dates, and indicators are deliberately kept separate.

| Indicator and source | Reported quantity | Scope and period | Interpretive boundary |
|---|---|---|---|
| GenAI occupational exposure [74] | About 25% exposed; 3.3% highest category | Global employment; 2025 methodology | Potential task exposure; not measured displacement |
| Care-related labour-force non-participation [73] | 748 million persons, including 708 million women | Age 15+; global estimates for 2023 | Reasons for non-participation; not total unpaid care hours |
| Informal employment [79] | 57.9% of employed people | Global estimate for 2025 | Employment status; not direct evidence of AI impacts |
| Algorithmic management [140] | 90% US; 79% mean across surveyed European countries | Employer/manager survey in six countries; published 2025 | At least one tool; not worker-level AI adoption or prevalence worldwide |
| Occupational segregation [78] | 29% of female-dominated versus 16% of male-dominated occupations exposed | ILO analysis of occupations; 2026 | Occupation-level proportions; not proportions of all women or men losing jobs |

Michaels studied industrial robotics across countries and industries and found productivity effects alongside a different pattern of aggregate employment results [50]. These results should not be combined as if they measured one universal employment effect of automation: identification strategies, units of analysis, sectors, technologies, and comparison periods differ. Generalizing either study to AI-directed material production in 2026 would require new data.

The International Labour Organization's 2025 refined index estimated that approximately one in four workers worldwide occupied jobs with some generative-AI exposure; 3.3% of global employment was classified within the highest exposure category [74]. The index estimates a capability-related exposure of occupational tasks, not observed adoption, job elimination, realized productivity, or future employment levels. A 2026 ILO methodological brief explicitly emphasizes the limitations of exposure measures as employment forecasts [80]. The exposure of a clerical task to a language model and the readiness of an autonomous industrial robot to replace a hazardous physical operation are different technical propositions; the paper does not use the former as a proxy for the latter.

Table 17 brings together indicators whose denominators and interpretations cannot be interchanged. Occupational exposure measures potential task change; care-related non-participation documents a structural condition of labour-market access; informality records differing relations to protection and regulation; algorithmic-management surveys describe particular kinds of software and surveyed workplaces. None independently identifies a causal consequence of autonomous AI production. This heterogeneity is informative: the social consequences of automation cannot be evaluated through a single series on jobs or installed equipment.

## 12.3 Maintenance Work and the Redistribution of Productive Responsibility

Sections 7 and 8 distinguished task execution from the continuing capacity to repair, replenish, diagnose, and reproduce productive infrastructure. This distinction creates a corresponding question about living labour. A robot that performs a repetitive assembly movement may still rely on technicians who maintain actuators and sensors, production engineers who validate process changes, workers who check component quality, software specialists who investigate failures, and upstream personnel who manufacture replacements. Some functions may themselves be partly automated, but demonstrable continuity still depends on the actual allocation of residual responsibilities and the resources available to discharge them.

It is therefore necessary to examine *maintenance labour* as a productive relation, rather than treat its existence as an automatic refutation of automation. The empirical tasks include identifying failure

DRAFT

frequencies, mean time to repair, procedures for detecting nonconforming outputs, time spent validating autonomous decisions, skills needed to recover from unexpected states, and the distribution of authority to stop unsafe operation. These activities differ from both routine production and unpaid domestic maintenance, although they can interact through working schedules and care responsibilities. A high rate of successful automated task execution may coincide with scarce diagnostic specialists or increasing on-call work. Alternatively, automation may eliminate hazardous exposure and improve maintainability. Both outcomes remain possible.

The skills required for technical renewal also emerge from institutions of training and apprenticeship. An automated facility may rely on accumulated knowledge transmitted through experienced personnel even if fewer workers participate in the immediately observable operation. The ILO's 2026 study of changing workplace skills reports changing demands for higher-order cognitive, socioemotional, digital, and AI-related skills, together with emerging specialized occupations concerned with AI development and maintenance [76]. These descriptions motivate investigation; they do not establish that all displaced workers can enter new positions, that training costs are equally borne, or that every region benefits. Transferability depends on credentials, time, finance, geographic mobility, accessibility, and the institutions through which qualifications are recognized.

## 12.4 Algorithmic Coordination and the Conditions of Working Life

Autonomous production can change who has the practical authority to organize work. Allocation of tasks, evaluation of output, scheduling of personnel, and monitoring of production may be conducted partly through automated systems. Such software can coordinate complex facilities, but it can also change visibility and control over workers' activities. The institutional questions concern the scope of monitoring, the contestability of decisions, responsibility for errors, opportunities for worker participation, and whether time saved by automation is redistributed as reduced working time, changed workloads, or higher production targets.

The OECD's 2025 study of algorithmic management surveyed managers in France, Germany, Italy, Japan, Spain, and the United States. It reported that 90% of surveyed US managers said their firms had adopted at least one of the defined tools to instruct, monitor, or evaluate workers; the corresponding average across the four European countries surveyed was 79% [140]. These are survey-based measures of algorithmic management within particular countries and a broad software category; they should not be read as evidence that 90% of US workers are supervised by autonomous AI agents, or that all such systems produce adverse consequences. The report also discusses managers' perceived decision-quality improvements alongside concerns involving accountability, interpretability, and worker health.

From a historical-materialist perspective, technologies of coordination modify the organization of the labour process and the distribution of discretion within production relations. Nevertheless, a particular technical architecture does not uniquely determine its social implementation. Worker consultation, collective representation, labour inspection, occupational safety standards, access to explanations, and genuinely exercisable appeal channels can alter the institutionally feasible organization of work. A comparative study must also address selection: firms adopting management software may differ in scale, labour composition, managerial strategy, or regulatory environment before adoption. Correlation between software use and reported stress or productivity is insufficient by itself to establish a causal effect.

## 12.5 Care, Subsistence, and the Renewal of Human Capacities

Social reproduction places the transformation of paid work within a larger material setting. Daily and intergenerational reproduction of human capacities requires time, food, housing, healthcare, education, relationships, care, and public and ecological infrastructures. These conditions do not follow mechanically from the output of a factory. They are organized through differentiated combinations of households, public institutions, communities, firms, and markets. Increasing throughput per labour-hour might expand the potential availability of goods and services while simultaneously reducing earnings for some workers, demanding more time for retraining, or placing greater costs on households. Whether productive gains

DRAFT

improve life conditions depends on real income, distribution, public provision, hours, accessibility, and the quality of services.

The ILO's 2024 estimates found that 748 million persons aged 15 years and above were outside the labour force in 2023 because of unpaid care responsibilities, including 708 million women and 40 million men, based on data from 125 countries and imputation for countries with missing information [73]. This is a modeled estimate of reported reasons for labour-force non-participation, *not* a headcount of all people performing care, an estimate of unpaid care hours, or evidence about AI adoption. It nonetheless demonstrates that the economic interpretation of automation is incomplete when capacity to undertake paid work is considered without the time and care conditions that make participation feasible. Care-dependent households may benefit from lower prices or improved public services, yet they may also bear costs when employment becomes unstable or work schedules less predictable.

Different populations experience such effects unevenly. In its March 2026 analysis, the ILO reported that female-dominated occupations were more likely to be exposed to generative AI than male-dominated occupations in the occupational sample it examined (29% compared with 16%) [78]. This reflects patterns of occupational segregation, not a universal causal effect of adoption on women or a prediction of individual job loss. Further differentiation is necessary by contract type, disability, age, migration status, household composition, region, and the capacity to access training and social protection. These variables are relevant to understanding how historically existing inequalities can condition transitions, without treating any demographic group as internally uniform.

A reproductive assessment would consequently ask who receives time freed by automation, whether families gain or lose reliable income, how educational and care institutions adjust, and whether displaced work is transferred to unpaid household members. Improved efficiency inside a production boundary may be socially costly if it relies on unstable schedules or if training and maintenance burdens are displaced to workers without compensation. Equally, accessible training and reliable care services could enable technical change to widen feasible participation. Such claims should be measured in physical time, income, service availability, occupational risk, and independently observed wellbeing, preserving distinct units rather than summing them into an ungrounded welfare index.

## 12.6 International Supply Relations and Differentiated Labour Visibility

The relational representation of productive dependence must include labour performed beyond the focal establishment. Hardware fabrication, mineral processing, logistics, data preparation, content evaluation, and digital services may involve workers employed under different contracts and legal regimes. Industrial automation could shift labour requirements toward other regions or sectors, or it could improve the productive opportunities of particular suppliers. To identify these changes, researchers must observe relevant upstream and downstream activities and distinguish employment opportunities created by greater demand from changes in conditions or control within existing workplaces.

The ILO's 2025 survey of 1,153 web-based platform workers in 21 Latin American and Caribbean countries reports conditions including supplementary-income reliance and incomplete social protection, within its regional and survey population [75]. Such evidence illustrates possible institutional forms of digitally mediated labour but does not represent all data work, industrial supply chains, or autonomous-production facilities. Likewise, the ILO's 2026 analysis across 135 countries found that digital infrastructure and national task structures condition the balance between exposure to generative AI and potential augmentation gains [77]. These findings argue against transferring the labour experience of a highly capitalized automated facility directly to every upstream supplier or labour market.

Environmental work also enters this accounting. Waste handling, remediation, ecological monitoring, and the maintenance of water and energy systems can support productive continuation while being omitted from narrow labour-productivity accounts. Their technical quality and ecological effects require independent evaluation. A proposal to increase the share of workers engaged in environmental tasks would have to ask what those tasks accomplish, what risks they create, how workers are compensated, and whether preventive changes could avoid the need for subsequent repair. Labour hours devoted to

DRAFT

remediation cannot themselves establish successful ecological recovery.

## 12.7 An Accounting Separation of Task Intensity and Paid Work

A minimal formal accounting identity helps prevent unwarranted deductions from improved task efficiency. Fix a production boundary $B$, two comparable periods $t \in \{0, 1\}$, and product categories $k = 1, \ldots, K$ with consistent unit definitions and quality requirements. Let $Q_{k,t}$ denote conforming output units, $a_{k,t}$ directly attributable *paid labour-hours per output unit*, $M_t$ paid infrastructure-maintenance hours not already counted in $a_{k,t}$, and $V_t$ paid coordination, validation, and recovery hours not already counted in either category. Define

$$H^{\text{paid}}_{B,t} = \sum_{k=1}^{K} Q_{k,t} a_{k,t} + M_t + V_t, \qquad [H^{\text{paid}}_{B,t}] = \text{person}\cdot\text{hours}. \tag{31}$$

The addends must be mutually exclusive, observed over a common horizon, and measured for the same production boundary. This is bookkeeping, not a production function or causal law. In particular, separately accounting for maintenance avoids hiding the labour needed to make an automated plant work. Where upstream labour is included, its scope and any double counting must be declared.

The exact change between the two periods can be decomposed using period-0 unit requirements as a reference:

$$\begin{aligned} H^{\text{paid}}_{B,1} - H^{\text{paid}}_{B,0} = {} & \sum_{k=1}^{K} (Q_{k,1} - Q_{k,0}) a_{k,0} + \sum_{k=1}^{K} Q_{k,1} (a_{k,1} - a_{k,0}) \\ & + (M_1 - M_0) + (V_1 - V_0). \end{aligned} \tag{32}$$

The first term accounts for output-volume differences at reference task intensity; the second records differences in direct paid hours per unit evaluated at period-1 quantities. The remaining terms record changes in separately measured maintenance and coordination. An alternative ordering yields a different two-factor decomposition but the same total change. No individual term identifies the causal effect of AI without a valid counterfactual and an account of other concurrent changes.

For illustration, consider a hypothetical single-product facility that initially produces 100 units using 2 paid direct hours per unit, 20 maintenance hours, and 10 validation hours, giving 230 paid person-hours. In a later period, output rises to 150 units while direct hours per unit fall to 1; maintenance rises to 40 hours and validation to 30 hours. The later total is 220 person-hours. The 50-unit output increase contributes $+100$ hours at the initial direct-hour coefficient, the intensity change contributes $-150$ hours, and maintenance plus validation add $+40$ hours, yielding a net change of $-10$ hours. Direct paid labour per unit has halved, yet total paid hours have fallen by only 10. Other parameter combinations could produce increased total hours or unchanged working time. These numbers are constructed solely to check and communicate the identity.

Converting paid person-hours into employed persons further requires information on hours per worker, turnover, labour contracts, and multiple jobholding. It cannot be inferred from Eq. (32) alone. Likewise, household care time is not an unmeasured remainder to be inserted as a residual in the paid-hours equation; it requires a separate observation scheme. A linked social-reproduction record might track hours of paid work, unpaid care, access to education and healthcare, occupational risks, real disposable income, and participation rights. These are heterogeneous variables with different ethical and evidential status, just as the productive-outcome tuple in Eq. (28) has no intrinsic common unit.

## 12.8 Comparative Research on the Continuation of Human Productive Capacities

A useful study would follow comparable facilities before and after the deployment of a defined AI-mediated production function, collecting task-level logs and linked, privacy-protective records of working hours, maintenance requests, qualifications, wages, contract stability, worker consultation, and occupational health. It would separately measure upstream labour changes and the provision of care, training,

DRAFT

and public services in the relevant labour markets. A second design would compare otherwise similar facilities operating under different worker-participation, training-finance, or working-time institutions. Neither comparison would be interpreted causally without considering adoption selection, product demand, pre-existing trends, outsourcing, business-cycle effects, and changes in workforce composition.

The central hypotheses are deliberately conditional. Where automation reduces direct task intensity, total paid work may fall, remain stable, or rise according to output expansion and residual labour requirements. The reproductive consequences may diverge even where total hours are similar, because income, contract security, scheduling, access to training, and care needs differ. The same machinery may be accompanied by different labour outcomes under different ownership and governance arrangements; differences must be measured rather than imputed from institutional labels. Finally, a more complete accounting of upstream, household, and environmental work may reverse a conclusion obtained from factory-level productivity alone. Evidence showing that broadening the boundary makes little difference, or that automated coordination improves job quality without increasing hidden labour burdens, would appropriately narrow these hypotheses.

Within the present paper, these are proposed tests, not observed effects of a new facility. The original analytical contribution is to connect the boundary-sensitive reproduction of production infrastructure to the socially situated reproduction of human capacities, while specifying the evidence needed at each interface. The link remains open to revision where existing labour economics, time-use research, and social accounting already explain the data adequately. The historical-materialist implication is that AI-mediated productive autonomy alters a set of socially organized possibilities whose distribution and continuation depend on labour institutions, physical resources, and the relations through which technological and human capacities are reproduced. The next section begins the engineering review by examining the demonstrated capabilities of industrial robotics; its technical performance claims must be kept separate from the social consequences analyzed here.

# 13 The Demonstrated Capabilities of Industrial Robotics

The preceding sections established how productive autonomy depends upon boundary selection, recursive infrastructural requirements, access to productive bottlenecks, and the social reproduction of labouring capacities. This section begins the engineering investigation by asking what industrial robotic systems have demonstrably achieved, within which operating conditions, and with what continuing reliance on human and infrastructural support. The research object is the *evidenced capacity of robot-mediated industrial tasks and integrated production cells*. It includes fixed manipulators, human–robot collaborative applications, mobile industrial platforms, and selected learning-based control policies. Their capabilities are compared on declared tasks rather than treated as interchangeable expressions of a single autonomy level. This analysis prepares the subsequent examination of autonomous laboratories, adaptive production control, computing infrastructure, energy and materials, and maintenance. Throughout, a robot's completion of a bounded operation is distinguished from the autonomous reproduction of the conditions that make the operation possible.

## 13.1 Installed Industrial Robots and the Meaning of Deployment

Industrial robots have become established productive equipment in automotive, electronics, metalworking, material handling, and other manufacturing activities. The International Federation of Robotics (IFR) reported approximately 603,000 industrial robots installed worldwide during 2025 and an operational stock of 5.079 million units at year-end, in its *World Robotics 2026* release of 24 September 2026 [69]. These figures measure different quantities: annual installations are a flow of equipment deployments; operational stock is the estimated population of industrial robots operating at a specified point in time. Neither series counts fully autonomous factories, and neither directly records the rate of production without human support, the quality of maintenance, embodied resource demands, or the ownership of resulting outputs. The installed population provides strong evidence that robot-mediated operations are

DRAFT

industrially consequential, but little by itself about the reproductive autonomy developed in the preceding formal sections.

The unit of evaluation consequently matters. A manufacturer can use a robot arm for positioning, welding, coating, packaging, loading, or inspection while the surrounding production system depends on fixtures, material delivery, programmable logic controllers, quality supervision, human escalation, and component replacement. Even where a robot operates unattended during a shift, its task autonomy can rely on extensive preparation outside that shift. Equating a deployed arm with an independently functioning factory would omit the productive relations through which workpieces, commands, tooling, spare parts, skilled services, and permissible operating conditions reach the robot. This is the industrial interpretation of the boundary-sensitive autonomy distinction established earlier in the manuscript.

A careful comparison needs to distinguish *equipment availability*, *cell-level task performance*, and *plant-level output*. Those indicators have different denominators: an installed arm can be available yet idle for lack of upstream inputs; a successful cell cycle can fail downstream quality requirements; and a factory can deliver conforming output while relying on frequent manual recovery. The corresponding evidence should record the functional boundary, specified product, runtime, number of cycles, intervention rules, and measurement method. Aggregate deployment statistics should be used as evidence of diffusion, not as substitute performance testing.

## 13.2 Manipulation Performance Requires Task-Specific Metrology

The engineering maturity of industrial manipulators is partly attributable to stable, precisely specified tasks and performance-test conventions. ISO 9283:1998, confirmed as current by ISO, specifies criteria and methods for pose and path accuracy, repeatability, stabilization, overshoot, velocity behaviour, and related manipulator characteristics [81]. Such measures are indispensable for comparisons of controlled motion, yet pose repeatability is distinct from absolute positioning accuracy and both differ from the probability of producing a conforming assembly. A repeatable end-effector can repeatedly miss a poorly registered target; a calibrated position can still yield a failed fastening operation because of compliance, contact forces, fastener geometry, tool wear, or workpiece variation.

An engineering assessment should therefore record joint configuration, end-effector tooling, calibration, reference-frame registration, payload, speed, environmental conditions, and the location of the metrology instrument. The testing conditions matter as much as the reported figure. NIST research on robot accuracy degradation emphasizes the relationship between tool-centre-point pose, loading, application accuracy, and verification of performance over time [157]. In a 2024 NIST study of continuous mobile-manipulator measurement, the authors identify further pose and registration uncertainties introduced when manipulation is coupled to a moving base [1]. These findings caution against transferring a fixed-base repeatability specification directly to a moving, contact-rich production task.

The materiality of performance measurement is also historically important. Precision depends on capital equipment, calibration artefacts, measurements, controlled fixtures, reference geometries, skilled integration, and continuing inspection. The apparent independence of robot motion is enabled by a large infrastructure of socially produced accuracy. A GR representation adds useful detail only if it tracks changing calibrations, qualified tooling, task-specific uncertainty, and the provenance of the measurements more effectively than conventional manufacturing engineering and quality systems already do.

## 13.3 Assembly Capability Depends on the Whole Work Cell

Structured handling and repetitive programmed trajectories are mature applications; flexible manipulation of variable parts and contact-rich assembly remains more difficult. The difference involves perception of uncertain objects, control of contact forces, gripping variability, recovery from partial failures, and coordination among sensors, machines, and human operators. NIST's completed *Performance Assessment Framework for Robotic Systems* project explicitly treated perception, mobility, dexterity, and safety as interacting determinants of autonomous assembly performance; its project description notes the dependence of successfully automated assembly on extensive fixturing and tooling in many cases

DRAFT

[132]. That statement concerns a research programme and engineering assessment problem, rather than a measured global percentage of assembly operations that are automatable.

A 2024 laboratory study involving the FANUC CRX-10iA/L collaborative robot assessed assembly of a product component and identified the significance of workspace layout, gripper design, programming, and small-part manipulation [8]. Its conclusions are properly interpreted as evidence for a particular laboratory configuration and task, not as proof of factory-wide unattended assembly. Similarly, a successful insertion or fastening trial requires identification of setup conditions, tolerance variation, number of trials, cycle-time and quality definitions, and the rules for restarting after a jam or misgrasp. Success after operator repositioning or manual tool replacement provides evidence of bounded task capability, but should not be counted as an uninterrupted autonomous cycle under a stricter boundary.

For an integrated cell, a useful evaluation unit is a *qualified task episode*: an interval beginning with specified inputs and system condition and ending with conforming output, declared nonconformance, or recorded intervention. The episode includes the behaviour of feeders, conveyors, sensing, fixturing, grippers, and supervisory software when these are necessary for success. Success statistics that omit aborted attempts, reset times, or human assistance cannot support claims of sustained productive autonomy. Conversely, highly structured repetitive production may achieve excellent industrial performance without general-purpose perception or learning. Technical specialization is an important accomplishment; its limits should not be mistaken for failure merely because it is not general intelligence.

## 13.4 Industrial Mobility Expands Reach but Introduces New Dependencies

Driverless industrial trucks, automated guided vehicles, and some autonomous mobile robots extend material handling beyond a fixed robot cell. They can move workpieces between stations, deliver supplies, and reposition tools. Their technical evidence nevertheless concerns particular navigational environments, payloads, traffic rules, floor conditions, localization infrastructure, and human–machine interfaces. ISO 3691-4:2023 provides published safety requirements and verification for driverless industrial trucks and their systems, including specified automatic and manual modes [91]. As of October 2026 a draft revision was under development; the published 2023 standard should therefore be distinguished from proposed future requirements.

A mobile manipulator combines the uncertainty of moving-base localization with arm kinematics and contact conditions. NIST's 2024 mobile-manipulator measurement report examines issues in system integration and measurement for simultaneous base–arm operation [1]. Such research demonstrates why manipulator specifications, base navigation benchmarks, and whole-system production success must be examined separately before their conjunction is claimed. The layout of a manufacturing plant, the availability of chargers and communication, the condition of traversable floors, and rules for pedestrian interaction are all parts of its productive dependency network.

Industrial mobility thus broadens feasible work arrangements without eliminating the institutionally organized surroundings of production. It may permit material transport without fixed conveyors, yet it can increase dependence on fleet scheduling, batteries, safety zoning, localization markers, and maintenance of shared traffic spaces. To compare alternatives, one must account for both removed and newly introduced infrastructures and measure their actual effects on throughput, downtime, exposure, and flexibility.

## 13.5 Collaborative Applications Require Integrated Safety Evidence

The term *collaborative robot* is frequently used in equipment marketing, but the safety of human–robot collaboration is an application-level question. ISO 10218-1:2025 sets robot-level safety requirements, while ISO 10218-2:2025 addresses industrial robot applications and cells, including integration and safeguarding [94, 95]. The revision of the ISO 10218 series also incorporated significant parts of the earlier ISO/TS 15066 guidance into the application-level requirements. Safe deployment depends on hazard identification and verified safeguards for particular tools, loads, velocities, workpieces, and interactions.

DRAFT

A robot with force-limited functions is not thereby proved safe for arbitrary close human operation, especially when a sharp tool or hazardous process changes the risk profile.

Experimental evidence confirms the need to separate safety, human workload, quality, and throughput. Fournier and colleagues studied 120 participants completing simple and complex assembly tasks with either a human collaborator or a robot collaborator. In their controlled study, human–cobot pairs showed some favourable outcomes in task success or workload-related comparisons, but completion took longer and involved more gestures within the experimental assembly tasks [42]. This is a meaningful human-factors experiment, not a causal estimate of general factory productivity across sectors. A different 2024 engineering study investigated dynamic human–robot task allocation with productivity and wellbeing objectives [18], but an optimization model or demonstrator should not be equated with a broadly validated labour outcome.

From the standpoint of productive reproduction, human participation can take several forms: planned shared tasks, instruction, exception handling, observation, quality assurance, preventive maintenance, and emergency intervention. These should not be collapsed into a binary category of either human-operated or autonomous. Comparative trials should report the intervention and failure-handling protocol, the worker's exposure and decision authority, and worker-centred outcomes separately from robot cycle time. The legitimate allocation of risks and responsibilities is a normative and institutional question; its resolution cannot be inferred solely from a lower task completion time.

## 13.6 Learning-Based Policies Extend Demonstrated Behaviours within Measured Domains

Recent robot-learning research has expanded the experimental scope of manipulation through demonstration datasets, cross-embodiment training, and adaptable control policies. The Open X-Embodiment Collaboration reported a pooled dataset covering 22 robot embodiments through a collaboration among 21 institutions, with 527 described skills and more than one million recorded robot trajectories. Its RT-X experiments found positive cross-embodiment transfer in tested settings [145]. The figures refer to the authors' particular collection and experimental evaluations, rather than counts of independently validated industrial deployments. These findings are important evidence that information from multiple platforms can assist learning on some downstream tasks. They do not establish that the resulting policies are certified for safety-critical manufacturing, guarantee generalization to arbitrary factory tooling, or can maintain the physical infrastructure required for continued operation.

For industrial relevance, a policy should be assessed against the production conditions it may encounter: new workpieces, surface variation, worn tools, illumination changes, imperfect calibration, sensor occlusion, payload changes, human presence, and out-of-distribution faults. Failure detection and controlled fallback are as important as average benchmark success. Moreover, learned-control systems introduce dependencies on annotated demonstrations, data quality, compute hardware, model versions, software provenance, and human judgement concerning deployment boundaries. A nominally general policy may therefore relocate some integration burdens rather than eliminate them. The manuscript's earlier dependency model is an appropriate place to represent that possibility, but prediction of its magnitude requires measured changes in engineering effort and operating performance.

The evidential hierarchy matters. Simulated successes, laboratory robot demonstrations, pilot manufacturing cells, and sustained operating deployment address different claims. Cross-platform learning is supported by controlled research results; indefinite autonomous manufacturing or fully autonomous equipment reproduction requires further evidence. No inference from the number of recorded demonstrations to a universal probability of task success is justified without task and data-distribution information.

## 13.7 Component Reliability Does Not Directly Establish Cell Reliability

For a specified industrial task and horizon, productive success may require a chain of conditions: workpiece provision, perception, grasping, compliant motion, verification, material transfer, and admissible

DRAFT

operation. A failure at one necessary interface can interrupt the episode even when every major component has performed well in isolation. This issue can be expressed using an elementary probability bound without assuming that failures are independent.

Let $E_i$ denote successful completion of necessary interface $i$ in a declared task episode, and define $F_i = E_i^{\mathsf{c}}$. If $S = \bigcap_{i=1}^{m} E_i$ is the event that all necessary interfaces succeed, then Boole's inequality gives

$$\max\left\{0, 1 - \sum_{i=1}^{m} \Pr(F_i)\right\} \leq \Pr(S) \leq \min_i \Pr(E_i). \tag{33}$$

The left inequality is a lower bound obtained by the union bound; the right follows because $S \subseteq E_i$ for every $i$. Both are classical probability relations, not new theorems. They apply to events defined on the same task-episode probability space under the stated operating conditions. They provide no estimate without observed interface-failure rates and remain compatible with arbitrary statistical dependence among failures. If success also depends on recovery within an allotted time, recovery is an additional explicitly defined event or is included in the episode-level $E_i$.

In a purely hypothetical three-interface cell with success probabilities $0.99$, $0.98$, and $0.97$ under one declared workload, Equation (33) constrains total success between $0.94$ and $0.97$. It would be unjustified to multiply the three probabilities without establishing an appropriate independence or conditional model. A production cell's actual throughput also depends on repair delays, rework, upstream and downstream buffers, and quality acceptance, which are not captured by episode success probability alone. This example is illustrative accounting and does not represent a measured robot installation.

The practical implication is that robotic reliability is a property of the integrated productive arrangement. Assessment should include uptime relative to scheduled demand, conforming units per hour with the product mix stated, fault recovery latency, maintenance labour-hours, spare-part usage, production lost to quality failures, and the number of successful episodes without operator intervention. It should retain the underlying event logs, task definitions, controller version, and sensor calibration. Such records create an empirical bridge from engineering performance to the reproductive-dependency network developed in Sections 10 and 11.

## 13.8 A Comparative Protocol for Industrial Autonomy Claims

Table 18 specifies the evidence boundaries required for distinct industrial robotics claims. It is a proposed assessment framework synthesized from the preceding sources, not a ready-made international performance standard. The purpose is to prevent a claim about one system boundary from being silently enlarged into a claim about another. It also identifies a practical counterfactual: whether a GR-informed trajectory and dependency record yields measurable improvements relative to competent industrial documentation, safety assurance, maintenance databases, and production-execution systems.

For a concrete empirical study, the protocol would compare a conventional programmed cell, an adaptively controlled cell, and, where technically appropriate, a learning-assisted configuration delivering a matched product specification. The evaluator should preregister quality tolerances, intervention accounting, product variation, production horizon, safety exclusions, expected asset utilization, and evidence needed for an acceptable comparator. Outcomes would include conforming throughput, failure and recovery distributions, engineering and maintenance labour, energy and consumable requirements, and effects on relevant production dependencies. Distinct institutional conditions and acquisition costs should be recorded rather than silently held constant where they are not. Results favouring conventional automation, showing no difference, or showing superior performance from a learning-assisted policy would all be informative.

## 13.9 Implications for the Reproduction of Productive Capacity

The engineering literature supports a bounded conclusion: industrial robotic technologies already perform economically significant tasks at scale, with well-developed measurement and safety disciplines for

DRAFT

Table 18: Evidence interfaces for claims concerning industrial robotic capabilities. The comparison is an analytical research protocol, not a ranking of technologies or certifications.

| Claim being made | Appropriate observations or records | Important limit of inference |
|---|---|---|
| Robot is deployed industrially | Installed units, application class, operational stock, location and counting convention | Does not measure end-to-end factory autonomy |
| Motion meets a specification | ISO 9283-style test protocol, load, pose and speed, calibration, measurement uncertainty | Does not establish task or assembly yield |
| Cell repeatedly completes a task | Defined workpieces, attempt denominator, conformity, cycle time, aborts and resets | A controlled trial may depend on fixtures and manual intervention |
| Mobile system supports production | Navigation and load tests, site maps, traffic conditions, localization accuracy, charge and recovery logs | Navigation success does not establish manipulation accuracy |
| Human–robot work is admissible | Cell risk assessment, verified safeguarding and human-factors evaluation | A robot-level property does not certify an application |
| Learning policy transfers | Held-out tasks, embodiment differences, failure rates, reproducible policy versions | Benchmark transfer does not certify industrial generality |
| Production can continue autonomously | Long-horizon operating logs, exceptions, maintenance inputs, replacement and supply dependencies | Does not independently establish ecological compatibility |

selected applications, while more flexible assembly, mobile manipulation, and cross-platform learning require context-specific evidence of integration, reliability, and intervention. This conclusion is compatible both with rapid further improvement and with persistent practical limits. The demonstrated ability to transform workpieces through mechanical action is not equivalent to the ability to replace the equipment, energy systems, knowledge infrastructures, and labouring capacities that enable that transformation.

In historical-materialist terms, a growing population of industrial machines embodies accumulated science, engineering labour, and investments in fixed capital. Their actual productive possibilities are constituted through historically organized access to equipment, inputs, qualified skills, safety requirements, and maintenance. Automation may change the distribution of manual work and productive authority, but the direction of those changes depends on the production relation and technical configuration. The specifically GR research opportunity is to make the trajectory of those changing requirements explicit and test whether doing so adds explanatory value beyond the robust measurement practices already available. Ecological compatibility remains an independent empirical and normative question requiring lifecycle material and energy accounts, site-dependent disturbance evidence, and defensible limits. The next section therefore examines autonomous laboratories, where the integration of planning, experimental execution, materials provisioning, and validation makes the boundary of technical autonomy especially visible.

# 14 The Material Boundaries of Autonomous Laboratories

The preceding section examined what industrial robots demonstrably accomplish within bounded manufacturing cells. This section investigates the related but distinct productive organization of autonomous laboratories: systems that select experiments, manipulate materials, obtain measurements, and use results to determine subsequent experimental actions. Its research object is the *material and epistemic conditions of experimental autonomy*, including the apparatus and social infrastructure required to reproduce experimental capacity. Laboratory autonomy is a particularly informative case for the present paper because an apparently seamless inference–experiment–inference loop depends on physically heterogeneous processes: preparing precursor stocks, handling samples, calibrating instruments, removing residues, maintaining software interfaces, identifying products, and deciding whether an observation war-

DRAFT

rants a scientific claim. The section examines published demonstrations within their declared boundaries, then develops an accounting representation and comparative protocol for testing productive reproduction beyond the individual experimental campaign.

## 14.1 Experimental Closure Depends on a Materially Organized Workflow

A self-driving laboratory typically joins an experiment-selection policy to apparatus for preparation, execution, characterization, and data analysis. For a trial index $k$, its simplified decision structure can be written as

$$u_k = \pi_k(H_k), \qquad x_{k+1} = F(x_k, u_k, w_k), \qquad y_k = G(x_{k+1}, \eta_k), \qquad H_{k+1} = (H_k, u_k, y_k), \quad (34)$$

where $H_k$ is the recorded experimental history, $u_k$ is a permissible experimental instruction, $x_k$ is the underlying experimental state, $w_k$ includes uncontrolled physical influences, and $y_k$ is an instrument record affected by measurement conditions $\eta_k$. These objects are typed by the experimental protocol; $u_k$ may comprise precursor identities, concentrations, furnace schedules, or process temperatures, while $y_k$ may contain spectra, diffractograms, images, or electrical measurements. Equation (34) is a schematic closed-loop representation, not a validated universal model of laboratory science.

An executable $u_k$ presupposes that the commanded operations are compatible with available reagents, containers, robot geometry, tool changeover, safety interlocks, and instrument schedules. Likewise, $y_k$ becomes usable evidence only through calibration, preprocessing, reference standards, and a defensible correspondence between signal and inferred property. Both conditions can fail even when the planning algorithm functions as specified. An experiment may be selected correctly but never completed because a solvent reservoir is depleted, a gripper loses a vial, or a balance is unavailable. Conversely, equipment may execute an instruction correctly while generating data that are too ambiguous to identify the desired material. Published reviews of self-driving laboratories emphasize the engineering work required to integrate instruments, control software, analytical routines, and laboratory operations [20, 188].

The boundary of autonomy consequently requires a description of *who or what performs each prerequisite operation*. A laboratory that selects and executes experiments for eight days may rely on people who previously supplied clean vessels, prepared stock solutions, validated calibration methods, maintained the instrument controllers, defined acceptable scientific targets, and later assessed the results. These contributions are constitutive elements of the productive process. Treating them as externalities of experimental success would reproduce the boundary errors identified in Sections 13 and the preceding dependency analyses.

## 14.2 Materialized Experimental Workflows in Demonstrated Systems

Several well-documented systems show different ways to combine experimental decision-making and physical operations. Steiner and colleagues implemented a modular chemical synthesis platform driven by a formal chemical programming language, demonstrating execution of specified organic syntheses [175]. Its principal contribution to this discussion is reproducible translation of procedural instructions into instrument operations; following a formally expressed recipe does not establish that the platform independently defines research objectives or replenishes its hardware. MacLeod and colleagues' *Ada* platform combined robotic thin-film preparation, optical and electrical characterization, and Bayesian search to optimize a defined materials objective [117]. That study demonstrates adaptive experimental optimization under a configured measurement and fabrication workflow. The reported thin-film metric should be interpreted within the study's characterization protocol rather than generalized to commercial device qualification.

Burger and colleagues reported a mobile robot that executed 688 photocatalysis experiments over eight days while searching a ten-variable formulation space. The selected mixtures were approximately six times more active than the study's initial formulations under the specified test conditions [16]. Here the robot's mobility enabled use of laboratory instruments designed for human operators. This is an

DRAFT

Table 19: Evidence boundaries in selected autonomous laboratory demonstrations. These descriptions summarize reported experiments, not independently repeated performance tests.

| Demonstration | Established within the study's scope | Further evidence required for reproductive autonomy |
|---|---|---|
| Steiner et al. (2019) | Modular execution of specified organic syntheses using programmable operations | Autonomous acquisition, maintenance, verification and adaptation of apparatus |
| MacLeod et al. (2020) | Closed-loop thin-film processing and optimization of a measured material property | Replication across feedstocks, devices, instruments and operating horizons |
| Burger et al. (2020) | 688 experiments in eight days within a photocatalyst-formulation search | Continued operation after replenishment, wear, recalibration and unexpected faults |
| Boiko et al. (2023) | Tool-assisted chemical planning and selected experimental operations | Independently measured robustness across laboratories, exceptional cases and controls |
| Szymanski et al. (2023; corrected 2026) | Robotic inorganic-powder synthesis and phase-informed optimization across 57 targets | Independent validation of identity and purity, supply renewal and long-run equipment continuity |

instructive architectural alternative to rebuilding every instrument as an autonomous module. Yet the comparison is confined to the declared photocatalytic performance measure and laboratory campaign: it is neither a claim of sixfold industrial energy efficiency nor evidence of autonomous reproduction of reagents, instruments, or research facilities.

The Coscientist system studied by Boiko and colleagues employed a large-language-model-driven tool architecture to search information, plan operations, generate laboratory commands, and undertake selected chemistry experiments, including reaction optimization [14]. Tool-mediated experiment design broadens the interfaces through which accumulated knowledge can guide experimental action. It also makes error propagation across documentation, code generation, command translation, safety permissions, and material execution an explicit concern. These studies collectively establish bounded and heterogeneous forms of experimental automation, adaptive optimization, and planning. Their different scientific objectives, intervention arrangements, and outcome measures preclude ranking them on an undifferentiated autonomy scale.

## 14.3 The Corrected A-Lab Case and the Warrant for Material Claims

The A-Lab supplies a particularly valuable case because it unites robotic powder handling, model-assisted synthesis planning, X-ray diffraction (XRD), computational materials information, and active learning [178]. Its published article was corrected on 19 January 2026 following concerns about the unambiguous identification of compounds from diffraction and interpretation of material novelty [179]. The corrected article reports that the system realized *36 compounds out of 57 target materials* during a 17-day campaign, with 353 experimental runs reported in the article. The authors' post-publication reanalysis found 36 supported results among 40 previously reported successes, with four results remaining inconclusive on the available diffraction evidence. Separately, one target was removed because it had mistakenly been included in the prediction-platform training data, accounting for the denominator revision. The correction also clarifies that "novelty" in the earlier presentation had been capable of suggesting novelty to science when some intended comparisons concerned novelty to a prediction platform. Thus the older figures of 41 successes among 58 targets should not be reused as current results.

The distinction is substantive rather than merely editorial. Detecting a signal compatible with a target crystalline phase is not necessarily equivalent to uniquely establishing its composition, phase purity, functional properties, or scientific novelty. The corrected article explicitly notes that some samples contained byproducts and that XRD-only evidence can leave ambiguity among candidate phases [178, 179]. Scientific judgments can require independent compositional analysis, additional diffraction or spectroscopy, repeat synthesis, and comparison with prior material records. The appropriate evidential threshold de-

DRAFT

pends on the claim. Identification of a major phase, demonstration of phase purity, and evidence that a composition has never been prepared before are different research propositions.

This episode should not be flattened into either a categorical rejection or a universal endorsement of autonomous chemistry. It demonstrates substantial integration of material handling, planning, and data processing during a bounded campaign. It also exposes a consequential weakness at the interface between an automated classification and the stronger proposition asserted from that classification. The system's ability to proceed through the experiment loop is an engineering achievement; the epistemic legitimacy of the resulting material claim requires scrutiny of its evidence and alternatives. A mature experiment-inference architecture should retain ambiguous cases as unresolved, preserve the basis of reclassification, and make independently reviewable corrections possible. The formal representation of these conditions is developed below.

## 14.4 Measurement Traceability and the Possibility of Independent Correction

Laboratory evidence is produced through chains of physical transformations and interpretations. A characterization result depends on instrument condition, calibration and reference standards, sample preparation, data acquisition, signal corrections, model assumptions, and expert or automated interpretation. ISO/IEC 17025:2017 specifies requirements for the competence, impartiality, and consistent operation of testing and calibration laboratories; the standard provides a relevant methodological benchmark but does not confer accreditation on any autonomous platform merely because it is cited [96]. NIST additionally specifies that metrological traceability is a property of a particular measurement result, supported by a documented calibration chain and uncertainty, not a property automatically inherited from a laboratory or instrument. Traceability alone does not establish that uncertainty is sufficiently small for the scientific decision at issue [156].

Consider an inferred scientific proposition $c$ obtained from a recorded signal $y$. If two experimentally admissible states produce observationally indistinguishable outputs, the inverse problem may be underidentified. A simple statement exposes this limitation without presupposing a particular laboratory instrument.

**Proposition 6** (Non-identification from an observationally indistinguishable pair)**.** *Let $\mathcal{X}$ be a set of physically admissible sample states and $g : \mathcal{X} \to \mathcal{Y}$ a deterministic measurement mapping. If there exist $x_1, x_2 \in \mathcal{X}$ with $x_1 \neq x_2$ and $g(x_1) = g(x_2)$, no estimator using only the single recorded value $g(x)$ can correctly distinguish $x_1$ from $x_2$ for both alternatives.*

*Proof.* For both states the estimator receives the same input $y = g(x_1) = g(x_2)$ and therefore returns the same output. It cannot return two distinct state identities from that common input while being correct for both. □

This elementary observation is not a novel theorem of spectroscopy. It states the assumption that makes independent discriminating measurements necessary. In a noisy stochastic setting, even unequal likelihood functions need not provide practically adequate discrimination at the available signal-to-noise ratio. The proposition is used to clarify the evidential burden, not to assert that particular A-Lab diffraction observations were mathematically identical. Independent assays, additional operating conditions, or repeat samples can sometimes distinguish candidate explanations, provided their assumptions and errors are recorded. The GR representation proposed earlier is useful only if it traces the dependencies from each scientific proposition to its material sample, instruments, transformations, models, and human decisions more effectively than competent laboratory information systems already do.

## 14.5 Consumables, Calibration, and the Reproduction of Experimental Capacity

Unlike a purely computational search, physical experimentation transforms inventories. Each run may consume reagents, substrates, disposable pipette tips, gases, sample holders, cleaning agents, and instrument time. It may leave waste and impose wear, thermal cycling, contamination, or calibration drift.

DRAFT

Many resources are heterogeneous and cannot responsibly be combined into an undifferentiated scalar "material input". For a given material species or separately measured inventory class $i$, a discrete balance can be expressed as

$$m_{i,k+1} = m_{i,k} + r_{i,k} - c_{i,k} - \ell_{i,k} + \rho_{i,k}, \qquad i \in \mathcal{I}, \tag{35}$$

where $m_{i,k}$ and $r_{i,k}$ are available inventory and replenishment, $c_{i,k}$ is experimentally consumed material, $\ell_{i,k}$ represents other losses, and $\rho_{i,k}$ represents recoverable material returned to the specified inventory. Terms must use compatible units for each $i$, and recoveries must not be counted twice. The representation can be extended to batch expiry, reagent quality, hazardous-waste capacity, and material substitution but does not itself predict those conditions.

Experimental continuity also depends on the condition of fixed assets. A robotic dispenser's metrological validity can change before mechanical breakdown. An instrument may operate but deliver systematically biased values following an unrecognized calibration change. Therefore the readiness of a laboratory at the next decision cycle includes both available materials and a verified capability to produce admissible measurements. Denote by $\mathcal{R}_k$ the set of experiments executable with currently qualified apparatus, inventories, procedural rights, and adequate safety margins. Successful completion of experiment $k$ does not imply that the next planned experiment belongs to $\mathcal{R}_{k+1}$. Reproductive laboratory autonomy requires specified conditions under which the relevant readiness set can be replenished or restored for subsequent campaigns.

A cost account likewise needs to separate variable supplies, fixed assets, qualified human support, and validation. For a specified evaluation period $H$, write

$$C_H = C_H^{\text{cons}} + C_H^{\text{energy}} + C_H^{\text{assets}} + C_H^{\text{maint}} + C_H^{\text{validation}} + C_H^{\text{human}} + C_H^{\text{waste}}, \tag{36}$$

where terms are monetary expenses under a declared accounting convention and currency. Separately report physical stocks and emissions rather than adding heterogeneous physical quantities into $C_H$. Asset costs require an explicit depreciation or replacement convention; human support includes qualified preparation, calibration, escalation, and interpretation. This account is a measurement proposal, not a claim that all expense categories are observable in published papers. A fair comparison with manual or conventional robotic laboratory practices requires matching product definitions, validation thresholds, experiment durations, and treatment of failed or repeated results.

### 14.6 Comparing Scientific Throughput with Validated Discoveries

A frequently reported output is the number of experiments performed, sometimes supplemented by the rate of objective improvement. Those measures are informative, but they answer different questions from the rate of reliable and independently confirmed discoveries. An integrated assessment should distinguish instructions generated, operations completed, analyzable records, claims passing predefined acceptance criteria, results independently replicated, and outcomes transferred beyond the original apparatus. Its denominator should be declared for each measure. Retrospective removal of failed or ambiguous experiments from denominators would overstate effective throughput.

As an explicitly hypothetical illustration, suppose an automated materials campaign schedules 80 trial conditions; 60 runs reach the prescribed completion state; 54 produce usable primary instrument records; 44 meet a preregistered candidate-screening criterion; and 40 survive an independent confirmation procedure. These are *constructed numbers*, not evidence from A-Lab or any published system. Their different fractions, $60/80$, $54/80$, $44/80$, and $40/80$, correspond to task execution, instrument yield, provisional identification, and validation. Reporting only 60 completed runs or only 44 screened candidates omits important evidence about the relationship between automation and scientific warrant. Under an alternative policy the number of completed operations could fall while the number of independently confirmed findings increases. Such a result would challenge an evaluation that assigns priority to nominal throughput alone.

Volk and Abolhasani emphasize the need for meaningful performance metrics across different self-driving laboratory architectures, while Canty and Abolhasani's 2026 review foregrounds scalability, gen-

DRAFT

Table 20: Prospective evidence grades for autonomous laboratory claims. The listed evidence is not presumed available for the published demonstrations.

| Claim being evaluated | Minimum relevant evidence | Insufficient inference |
|---|---|---|
| Experimental execution | Prespecified protocol, logs of completions, interventions and failed runs | A planned experiment counts as a physically completed one |
| Adaptive optimization | Defined objective, baseline policy, comparable experimental budget, repeated evaluation | More trials necessarily produce better knowledge |
| Product or phase identity | Measurement uncertainty, discriminating characterization and prior-art comparison | One classification establishes purity and novelty |
| Reproducible result | Independent replication, provenance and measurement comparability | A single-run instrument trace certifies reliability |
| Long-run laboratory continuity | Supplies, calibration, component failures, qualified repair, changeovers and recovery over stated horizon | Unattended operation implies infrastructure reproduction |
| Ecological compatibility | Lifecycle flows, energy and waste accounting, impacts, alternatives and ecological constraints | Higher scientific throughput implies smaller total burden |

eralizability, and provenance-complete experimentation [188, 20]. The present paper adds the question of the *reproduction of the capacity to produce those measurements*. This adds a material accounting and historically institutional dimension to reproducibility, but must be evaluated against established provenance practices rather than proclaimed intrinsically superior to them.

## 14.7 Institutional Reproduction and the Circulation of Experimental Knowledge

The equipment and knowledge required for an autonomous laboratory arise through diverse productive relations. A university platform, a proprietary research facility, a shared public instrument center, and a commercially operated cloud laboratory may perform technically similar operations under very different regimes of access, pricing, intellectual property, repair rights, specialist labour, and data circulation. These institutional conditions can shape who can submit experiments, inspect protocols, validate claims, or reuse results. They may also affect replacement capacity when a supplier ends a product line or withholds an instrument interface. A seemingly generalizable laboratory protocol can become non-portable if dependent on undocumented firmware, tacit skills, or non-transferable access credentials.

Historically, the automation of discovery embodies accumulated public and private scientific labour in specialized fixed capital. The experimental action that appears local to the laboratory is connected to upstream material purification, instrument manufacture, calibration services, trained technicians, and scientific conventions of identification. The question of appropriation, introduced in Section 11, therefore extends beyond ownership of a successful sample to the conditions that make experimentation possible and the institutions that recognize its output. Scientific productivity, monetary return, research access, and the capacity to sustain future inquiry can diverge in either direction. This paper accordingly proposes tracing productive relationships without preassigning ecological merit to commercial, public, or commons-oriented organizational forms.

A useful comparative study would evaluate matched scientific tasks conducted through a conventional manually supervised workflow, a programmed robotic workflow, and a self-driving laboratory. It would log researcher's preparation and rescue work, reagent use and waste, capital and calibration demands, valid outcomes, independently reproduced results, and access or governance restrictions. The central counterfactual would ask whether automation reduces total resource and labour requirements *per independently validated scientific result* under equivalent epistemic standards. Competent human practice could outperform the autonomous workflow on a specific study; additional experimental coverage

DRAFT

could improve knowledge at a higher material cost; and platforms with similar nominal autonomy could differ markedly in recovery time after supply or instrument failures. Each finding would refine rather than refute the broader historical-materialist inquiry into productive reproduction.

### 14.8 Consequences for Autonomous Productive Reproduction

The engineering record warrants confidence that integrated robotic planning, execution, characterization, and adaptive experimentation can operate effectively within selected scientific domains. The corrected A-Lab experience simultaneously demonstrates why the transfer from instrument outputs to scientific propositions should remain open to independent scrutiny. A completed experiment is a physical event; an interpreted measurement is an inference; a validated scientific discovery requires a claim-specific evidential process. Reproducing the productive conditions of autonomous experimentation introduces further requirements involving materials, maintenance, qualified people, knowledge infrastructures, and institutional access.

In historical-materialist terms, autonomous laboratories reveal how collectively produced intellectual capacities become operative through machinery and scientific institutions while leaving their material reproduction distributed across wider production networks. The GR research programme can contribute by tracking the changing relations among these layers, but the practical value of this additional representation remains an empirical matter. Ecological compatibility cannot be deduced from laboratory autonomy or from the speed of experimental optimization. It requires independent investigation of energy, material circulation, waste, displaced burdens, and the ecological conditions of continued production. The next section accordingly turns from the laboratory as an integrated demonstration to the control architecture that enables adaptive productive operations across wider industrial systems.

## 15 The Control Architecture of Adaptive Production

The previous section showed how an autonomous laboratory coordinates planning, material handling, instrumentation, and interpretation within a bounded experimental campaign. This section investigates the control architecture by which industrial production systems adapt their actions to measurements, disturbances, changing orders, equipment condition, and constraints. Its purpose is to identify the physical and informational interfaces at which apparent autonomy depends upon sensing, mathematical models, computation, organizational authority, and human intervention. The discussion separates well-established control methods and manufacturing integration standards from prospective claims about AI-enabled production, and then asks what evidence is required before describing an integrated production system as adaptive, safe, reproductively autonomous, or ecologically compatible.

### 15.1 Adaptation Takes Place Across Distinct Control Responsibilities

An industrial production process involves multiple kinds of decisions with different response times and evidential demands. Low-level controllers regulate motion, temperature, pressure, or chemical composition over short intervals. Supervisory controllers coordinate machines, equipment states, and fault responses. Manufacturing operations systems schedule jobs, allocate materials, and manage maintenance, while enterprise systems set broader demand and financial requirements. The ISA-95/IEC 62264 framework describes relevant interfaces between manufacturing operations management and enterprise planning, without specifying that a complete factory must implement one particular controller [62, 99]. Its significance here is organizational: optimization performed at one level is conditioned by constraints, data definitions, and decision rights established elsewhere.

Consider an order scheduling system that increases a furnace's planned utilization. A local controller may still regulate furnace temperature within permissible limits, while a maintenance rule may prohibit another cycle before inspection. A supervisor might defer the order because a sensor's calibration has expired, and an authorized human might halt production after an environmental alarm. The production

DRAFT

Table 21: Distinct interfaces within an adaptive production architecture. The timescales are qualitative and depend on the process; the entries define verification tasks, not universal industrial practice.

| Interface | Principal decision or record | Evidence required for an integrated-system claim |
|---|---|---|
| Equipment feedback | Actuation from measured state | Sensor calibration, actuator limits, disturbance response |
| State estimation | Inference from partial observations | Observability assumptions, uncertainty bounds, delay logs |
| Supervisory control | Modes, interlocks, fallback | Hazard tests, transition records, independent protection |
| Production planning | Orders, capacity, inventory | Executability, materials availability, schedule feasibility |
| Digital representation | Predicted process or asset state | Synchronization latency, validation domain, model residuals |
| Institutional access | Permissions, accountability, review | Decision rights, operator authority, audit and remedy routes |
| Ecological evaluation | Exposure and cumulative burdens | Independent measurements, system boundary, causal design |

outcome cannot be attributed to the scheduling algorithm alone. Adaptation is an interaction among controllers and resources whose responsibilities may partly overlap, contradict one another, or be incompletely specified.

This layered picture should not be interpreted as a universal chronology from conventional automation to AI-driven autonomy. A plant may employ advanced process control with little machine learning, or use machine learning only for maintenance prediction. Conversely, a generative model may help formulate scheduling options while the final commands are executed by well-established deterministic controllers. Those arrangements differ in their dependence on model training, inference services, authorization rules, and expert judgement. The research problem is to identify which function actually changes when AI is introduced and what additional requirements make that change reliable.

## 15.2 Partial Observation Conditions the Feasibility of Feedback

Physical control begins with a mismatch between the actual productive state and the information available to the controller. In a discrete-time representation, let $x_k$ be a vector of physically typed states, $u_k$ the applied command, $w_k$ a disturbance, and $y_k$ the recorded measurement. A general model is

$$x_{k+1} = f(x_k, u_k, w_k), \qquad y_k = h(x_k) + v_k, \tag{37}$$

where $v_k$ denotes measurement error. Physical typing matters: temperatures, tool wear, reagent inventories, electrical power, and product quality cannot be added as interchangeable scalar quantities. A production controller may estimate the state from an observation history $H_k = (y_0, u_0, \ldots, y_k)$, giving $\hat{x}_k = \mathcal{E}_k(H_k)$, while maintaining a set $\mathcal{X}_k$ of states still consistent with the observations and assumptions.

This representation carries no guarantee that the relevant state is observable. A sensor may fail to distinguish surface contamination from acceptable variation, or a temperature reading may conceal a local hot spot. A controller may also act on measurements collected before a material or equipment state changed. Thus a state estimate must be associated with its measurement provenance, acquisition time, model version, uncertainty, and domain of validity. When such information is missing, a record of apparently well-controlled equipment can overstate the evidence for reliable production.

For a particular plant, the uncertainties include stochastic measurement noise, systematic bias, sensor drift, disturbances, unmodeled effects, and discrepancies between command and actuation. These should not be conflated. Statistical prediction intervals typically require assumptions about the sampled process and their stated coverage; worst-case bounds require credible restrictions on possible errors. Neither may

DRAFT

hold after a process change without validation. The distinction becomes consequential when a controller operates near temperature, pressure, collision, or emission limits.

## 15.3 Predictive Control Coordinates Performance and Constraints

Model predictive control (MPC) provides an established framework for selecting control actions while explicitly representing future constraints. At decision time $k$, a typical formulation selects a sequence $u_{k|k}, \ldots, u_{k+N-1|k}$ by solving

$$\begin{aligned} \min_{u_{k:k+N-1|k}} \quad & \sum_{j=0}^{N-1} \ell(x_{k+j|k}, u_{k+j|k}) + V_f(x_{k+N|k}) \\ \text{subject to} \quad & x_{k+j+1|k} = \hat{f}_k(x_{k+j|k}, u_{k+j|k}), \\ & (x_{k+j|k}, u_{k+j|k}) \in \mathcal{Z}_k, \qquad x_{k|k} = \hat{x}_k. \end{aligned} \tag{38}$$

Only the first commanded action is applied before the problem is updated with new observations. The objective terms must be given compatible units or explicitly normalized weights, and the feasible set $\mathcal{Z}_k$ must correspond to real equipment and process constraints. This standard receding-horizon idea has a mature literature, including conditions for stability, feasibility, and constraint satisfaction [126, 158].

The optimization problem alone is not a certification of plant safety. A solution can satisfy a model's constraints while violating physical limits when the model is inaccurate, a sensor is biased, computation misses a deadline, or an actuator does not execute the assumed command. Robust MPC addresses some such difficulties by considering admissible uncertainty sets; stochastic approaches instead place probabilistic conditions on selected outcomes. Their guarantees differ and depend on the relevant assumptions. The relationship between digital prediction and physical execution therefore has to be evaluated through bounded-error tests, disturbance campaigns, failure injection, and explicit computation-time measurements.

Adaptive production may also modify a model as operating conditions change. Updating the model can improve tracking when wear, feedstock properties, or environmental conditions evolve, but it can invalidate constraints proved for an earlier model. A defensible adaptive architecture should therefore specify what updates are allowed in service, what new validation is required, and how control authority reverts to a verified mode when new predictions cannot be trusted. In many settings, learning is most appropriately used to propose candidate actions or refine forecasts under a separately enforced operational envelope.

## 15.4 Digital Twins Require Verified Correspondence with Production Processes

Digital twins are frequently presented as a means of connecting simulation, physical assets, and operational decision-making. The ISO 23247 manufacturing framework provides general principles and a reference architecture for such representations [89, 90]. NIST research identifies interoperability and the quality of synchronization, models, and computations as continuing implementation problems, particularly in additive manufacturing [168, 116]. The existence of a corresponding virtual object does not demonstrate that the representation predicts a changing physical process with sufficient accuracy for control.

A scientifically useful distinction is among a static engineering model, a simulation calibrated for a defined operating regime, a monitoring representation updated from equipment observations, and a closed-loop digital twin whose predictions actually influence physical commands. The required evidence changes accordingly. A simulation may reproduce a historical temperature profile but fail to predict defects under a new material batch. A synchronized model may expose this discrepancy promptly while lacking the means to diagnose it. A closed-loop controller may require stricter bounds on model residuals, update delays, and permissible operational changes than either monitoring or retrospective simulation.

DRAFT

Let $\hat{y}_{k+1|k}$ be a digital model's prediction of a subsequent measured quantity. The residual

$$r_{k+1} = y_{k+1} - \hat{y}_{k+1|k} \tag{39}$$

can be examined by operating condition, equipment age, batch, or material supplier. Small residuals in an aggregate dataset do not establish dependable behaviour near safety boundaries or under rare disturbances. Residual evaluation should also distinguish the data used for model fitting from independent validation and record intervals for which observation is unavailable. The relevant standard of evidence is the intended use: a model adequate for monitoring may be inadequate for high-consequence actuation.

A further historical-materialist issue concerns the ownership and maintenance of these representations. Models depend on process knowledge, sensor histories, supplier information, engineering conventions, and labour performed across organizations. Proprietary interfaces can restrict who may verify, modify, or repair the model. Alternatively, documented interfaces and reproducible validation procedures may expand the range of qualified maintainers. Whether digital integration increases productive autonomy or creates a new infrastructural dependency depends on these actual arrangements, connecting the present section to the bottleneck and appropriation analyses in Sections 9 and 7.

## 15.5 Safety Guarantees Depend on Explicit Error and Intervention Bounds

Safety-related control introduces obligations that differ from optimizing throughput or average error. Published work on control barrier functions, for example, derives conditions under which a dynamical system can remain within an admissible region when specified mathematical and input constraints hold [5]. Industrial deployment additionally requires equipment-specific hazard analysis, independent protective functions where appropriate, validated fallback behaviour, and conformance assessment under applicable standards. A proof for a mathematical plant is not a proof that a physical production installation satisfies all these requirements.

The effect of model uncertainty can be expressed with a deliberately elementary proposition. Let $g(x, u) \leq 0$ express one scalar physical constraint for a next state, and let the controller predict $\hat{x}^+$. Suppose the actual next state $x^+$ satisfies a credible error bound $\|x^+ - \hat{x}^+\| \leq \varepsilon$, while the constraint function is $L_g$-Lipschitz in its state argument for the relevant domain and fixed $u$. The model-predicted action is accepted only if $g(\hat{x}^+, u) \leq -L_g\varepsilon$.

**Proposition 7** (Margin required by a bounded state-prediction error)**.** *Under the stated norm, domain, and error-bound assumptions, an action satisfying $g(\hat{x}^+, u) \leq -L_g\varepsilon$ also satisfies $g(x^+, u) \leq 0$.*

*Proof.* The Lipschitz assumption gives

$$g(x^+, u) \leq g(\hat{x}^+, u) + L_g\|x^+ - \hat{x}^+\| \leq -L_g\varepsilon + L_g\varepsilon = 0. \tag{40}$$

□

This is a direct application of a standard Lipschitz inequality, not an original control theorem. Its value here is to expose the evidential premises hidden by the phrase "constraint-aware AI." If no justified upper bound for prediction error is available, the displayed argument gives no safety guarantee. Moreover, satisfying a single next-step inequality need not ensure a recursively feasible policy over later steps. Execution delay, cyber compromise, sensor faults, actuator saturation, simultaneous constraints, and emergency transitions require additional models and tests. The numerical size and legitimacy of $g$, $L_g$, and $\varepsilon$ must come from application-specific evidence, not from a generic claim about artificial intelligence.

An illustrative temperature-control example makes the distinction concrete. If a maximum allowed process temperature is $200\,^{\circ}\mathrm{C}$, an appropriately validated one-step prediction-error bound is $2\,^{\circ}\mathrm{C}$, and the constraint is $g(T, u) = T - 200\,^{\circ}\mathrm{C}$, the controller should require a predicted temperature no higher than $198\,^{\circ}\mathrm{C}$ to invoke the proposition. The calculation changes immediately if the error bound is invalid or an omitted hotspot temperature, rather than the measured bulk temperature, is the actual safety-critical variable. The figures are constructed for illustration and do not constitute an industrial operating recommendation.

DRAFT

## 15.6 Production Scheduling Does Not Inherit Machine-Level Guarantees

A production system can comprise individually controlled machines while suffering from failures of coordination among them. Job scheduling allocates equipment time, inventories, transport, utilities, and people, often in the presence of uncertain order arrivals or processing times. It may be formulated as a discrete optimization problem, whereas machine controllers regulate continuous or hybrid dynamics. Combining them requires an interface translating scheduled jobs into executable instructions and translating equipment failures or measured shortages into feasible schedule revisions.

The constraints are often coupled. A locally optimal machine action may consume a component that a downstream line requires; an energy-saving schedule may create a new thermal peak or increase product scrap after restart; a rapid response to an order may conflict with maintenance windows. These conflicts can be analyzed without supposing that one central AI agent must resolve all decisions. Distributed optimization, rule-based supervisory control, human approval, and hierarchical scheduling may each be appropriate under particular conditions. What matters for comparison is demonstrable performance under the same orders, uncertainties, available resources, and safety requirements.

A key distinction concerns the interval between detecting a deviation and completing a corrective action. A scheduler that instantaneously replans in software may remain dependent on technician availability, spare parts, supplier lead times, or physical equipment clearance. Thus decision latency, actuation latency, and restoration latency are distinct. Evaluating only software response can exaggerate operational autonomy and obscure the reproduction conditions examined in Section 14.

## 15.7 Cybersecurity and Institutional Authority Constrain Control Access

Digital production networks bring new possibilities for observation and coordination, together with vulnerabilities introduced by connectivity and software dependencies. The ISA/IEC 62443 standards family addresses cybersecurity processes and requirements for industrial automation and control systems, including responsibilities shared among asset owners, integrators, suppliers, and service providers [100]. NIST's AI Risk Management Framework also provides a voluntary, cross-sector approach to governing AI risks, though it should not be mistaken for a plant-level safety certification [180]. These documents address different questions: industrial-control cybersecurity requirements and voluntary AI-risk management are not, by themselves, demonstrations of machine-safety conformity, ecological permission, or authority to restart a hazardous process. Both remain relevant to productive continuity because a controller's feasible action set is affected by authentication, updates, remote access, and responsibility for recovering from incidents.

An autonomous planner may have technical capacity to propose a command without possessing authority to issue it. The difference is observable in command signing, access-control logs, emergency-stop rights, permitted operating envelopes, and escalation procedures. These institutions can improve accountability while limiting immediate flexibility; they may also create new dependencies on vendors or central administrative services. Conversely, unrestricted command access can undermine both physical safety and the reliability of scientific evidence about what actually occurred in production. The same control network can therefore be studied as a site where productive force, institutional authority, and potential power over continuation intersect.

Ecological permissions need independent treatment. The optimization of electricity cost does not establish that electricity supply is ecologically acceptable, just as a process meeting technical temperature constraints does not demonstrate that its cumulative material emissions or supply-chain impacts are acceptable. Where ecological constraints are modeled, their construction requires independent ecological observations and normative or legal decisions concerning acceptable thresholds. The controller can then be evaluated for compliance with those specified conditions; it cannot itself establish their scientific or institutional legitimacy.

DRAFT

Table 22: Proposed comparative measurements for adaptive production. Each indicator has a distinct meaning and cannot be inferred from a controller's nominal autonomy.

| Claim evaluated | Measurement and comparison | Principal limitation or negative finding |
|---|---|---|
| Tracking performance | Physical error under matched disturbances | Average accuracy can conceal rare unsafe states |
| Production continuity | Valid output and recovery time by failure mode | Depends on repairs, consumables and staffing |
| Control robustness | Constraint violations with model and sensor faults | Requires justified fault and uncertainty coverage |
| Adaptive benefit | Out-of-sample production performance | Improvements may disappear after process change |
| Computational burden | Latency, electricity use and service dependence | Software costs can displace material burdens |
| Institutional reviewability | Logged commands, approvals and interventions | Traceability alone does not establish rightful authority |
| Ecological consequence | Independently measured lifecycle and local effects | Physical optimization does not establish ecological viability |

## 15.8 Comparative Tests of Adaptive Productive Capacity

The proposed research programme should compare competent alternatives rather than a poorly implemented conventional baseline against an idealized AI system. A suitable experimental design would use the same specified production task, equipment, safety envelope, evaluation horizon, and disturbance scenarios for: a well-configured conventional feedback and supervisory system; a predictive controller with a validated process model; and an AI-assisted adaptive controller subject to the same independent safety and access constraints. Disturbances could include sensor drift, changes in feedstock, late orders, equipment failure, communication loss, and variations in material availability. If an algorithm requires additional cloud computation, model retraining, or technician support, those inputs should be counted within the chosen system boundary.

Each comparison must specify which claims can be identified. A measured reduction in temperature error supports a statement about temperature regulation under the tested conditions. It does not by itself support a reduction in total waste, an improvement in worker conditions, or an inference that the system can maintain itself for a longer horizon. A system may outperform a baseline under normal operation but fail more severely after network loss. An advanced controller may improve product quality while consuming more computation or requiring more expert intervention. Null and adverse results are substantive outcomes because they reveal the conditions under which additional control complexity yields little productive benefit.

The most demanding evaluation examines the connection to productive reproduction. A controller that keeps equipment within operating limits can contribute to component longevity, yet it may also increase wear through frequent adjustments or depend on sensors that require intensive recalibration. Testing this relation requires longitudinal maintenance and replacement records, not only short campaigns reporting throughput. The proposed GR representation would contribute only if its event and dependency records permit better identification of such effects than competent existing manufacturing information systems and provenance practices. That comparative possibility remains a hypothesis.

## 15.9 Consequences for the Reproduction of Autonomous Production

The engineering literature establishes that feedback, predictive optimization, system identification, and digital representations can expand productive coordination under specified constraints. Their deployment also makes visible a network of supporting activities: measurement calibration, model maintenance, communications, software security, equipment servicing, scheduling, supply provision, and institutional oversight. Autonomous control is consequently a situated achievement of organized productive rela-

DRAFT

tions. Extending control sophistication may reduce some manual tasks while relocating dependence toward model providers, instrumentation, computation, or specialized maintenance; it may also improve efficiency or resilience under appropriate designs. The direction is an empirical question.

For historical materialism, the relevant transformation is the incorporation of collective scientific and technical capacities into the control of material production, mediated by ownership, access rights, skilled labour, and organized infrastructure. The formal distinctions developed here help locate the conditions under which those capacities are exercised and reproduced. They also impose a limit on any ecological inference: adaptation and control can enforce well-specified physical constraints, but the adequacy of ecological constraints, their cumulative effects, and their institutional legitimacy require additional inquiry. The following section therefore examines the computing and communication infrastructures on which advanced coordination itself depends.

# 16 The Computational Dependencies of Productive Autonomy

Section 15 distinguished physical feedback, supervisory decisions, predictive models, and the authority to modify productive operations. This section investigates the computational infrastructure upon which those arrangements depend. Its objective is to specify when processing, storage, network communication, software maintenance, and external services become conditions of material production; to distinguish performance within an automated plant from the reproduction of its computing capacity; and to identify empirical tests of alternative digital architectures. Computation is considered here as physically implemented and socially organized activity. The treatment draws upon established operational-technology engineering, networking, data-centre accounting, and provenance methods while asking what further information is required to analyze changing productive dependencies relationally. It does not infer comprehensive factory autonomy from a powerful model or a successful isolated inference.

## 16.1 Computational Tasks Have Different Relations to Physical Production

Production-computing requirements are heterogeneous. A motion-control loop can require bounded response latency, while demand forecasting may tolerate periodic recalculation. An image-inspection system can depend on local camera data, accelerator availability, and a calibrated classification threshold; enterprise scheduling may rely on remote application interfaces and inventory records. Training a new model and executing an already validated inference model have different requirements for computational throughput, data access, energy, and organizational review. Consequently, a single measure of processor capability, such as nominal floating-point operations per second, cannot establish the capacity of an integrated facility to maintain production.

An industrial computing architecture commonly spans sensors and controllers, industrial networks, local or edge servers, and enterprise or remote computing services. The enterprise-integration architecture discussed in Section 15 provides one basis for distinguishing responsibilities across these interfaces. IEC 62264-1:2013 describes interfaces between enterprise and manufacturing-operations functions; ISO/IEC 30141:2024 instead provides a general Internet-of-Things reference architecture [62, 98]. Neither reference architecture specifies the correctness of a particular production algorithm, proves the availability of an implemented interface, or establishes the authority to act on a model output. An automation architecture must be reconstructed from its actual task requirements, data flows, fallback procedures, and maintenance provisions rather than inferred from its adherence to a reference diagram.

A critical distinction concerns *computational autonomy*: the ability to execute a specified computational function within a declared boundary and service budget. It differs from *productive autonomy*, which additionally requires physically feasible operations, materials, equipment, safety procedures, and institutional authorization. A facility might retain basic robot control during a remote AI-service outage while losing high-level scheduling or inspection capability. Conversely, a site with highly capable local models may still depend on externally supplied software updates, accelerator replacements, or qualified

DRAFT

Table 23: Computational interfaces and the evidence needed to establish their contribution to productive continuity. The examples are analytical categories, not claims about a particular factory.

| Interface | Typical productive dependence | Evidence needed for a system-level claim |
|---|---|---|
| Equipment controllers | Stable closed-loop actuation | Timing traces, bounds on jitter, tested local fallback |
| On-site inference | Inspection or manipulation policies | Model version, validation set, accelerator load, accuracy shift |
| Industrial network | Timely sensor and command exchange | End-to-end delay, congestion, packet-loss and recovery tests |
| Site servers and storage | Orders, histories, configurations | Capacity, backup restoration, update integrity, recovery time |
| Remote computing services | Model training, planning, analytics | Service dependencies, outage tolerance, contractual and access rules |
| Hardware supply and repair | Replacement of computing equipment | Spare-part availability, repair lead times, component provenance |
| Ecological accounting | Indirect material and energy burdens | Metered power, equipment inventories, system boundary and lifecycle model |

security support. These are different dependency structures, rather than higher and lower positions on a universal autonomy scale.

## 16.2 The Physical Reproduction of Computing Equipment

Processors, accelerators, memory, storage devices, networking equipment, power conversion, and cooling are material components of computing capacity. Their operation also depends on software layers including firmware, operating systems, drivers, orchestration tools, and production-specific configuration. Replacement capacity requires access to compatible hardware, qualified integration work, test instruments, and sometimes licensing or vendor support. The semiconductor supply dependencies discussed previously therefore remain relevant even where the plant's day-to-day computation appears independent of external intervention.

An analysis of this dependence should distinguish at least three time horizons. During a production shift, sufficient installed computing capacity and bounded latency may be decisive. Over a maintenance interval, secure updates, component repair, storage restoration, and diagnostics become critical. Over a longer renewal interval, equipment availability, compatibility migration, model retraining, energy infrastructure, and the continued existence of suppliers condition technological reproduction. A plant that can run an inspection model for one year with unchanged hardware has not thereby demonstrated that it can fabricate, repair, or replace that hardware without external support.

Hardware and software dependencies also exhibit complementarity: an available accelerator may be unusable with an obsolete driver, and a preserved model checkpoint may not run correctly on a changed processor or instruction set. Consequently, inventories need functional compatibility records, not merely counts of devices. The relational representation of Section 6 can associate a productive task with a set of jointly necessary computational artifacts and with feasible alternative paths. For these records to support comparison, each path requires test evidence for interchangeability, restoration time, and any altered quality or safety characteristics.

## 16.3 The Placement of Computation Changes Timing and Dependency Boundaries

Placing a computational function on a device, at the site edge, or in a remote data centre changes its dependence on communication paths, installed capacity, external power, provider access, and maintenance. Edge computing has long been studied as a way to address latency and bandwidth limitations of remote

DRAFT

services [164]. The relevant industrial decision, however, is conditional on a particular control or production task. Neither local execution nor remote execution is generally superior across all requirements. For instance, cloud infrastructure may provide redundant model-serving capacity and specialized maintenance, while locally controlled equipment may continue operating during a wide-area communication failure.

For a single sequential decision path, define the task-response time (in seconds) as

$$T_{\text{response}} = T_{\text{acquire}} + T_{\text{transmit}} + T_{\text{queue}} + T_{\text{compute}} + T_{\text{act}}. \tag{41}$$

The terms represent sensor acquisition, data transport, service waiting, computation, and actuation respectively, each measured from a declared start and completion event. They should be estimated from relevant trace distributions rather than a vendor's nominal round-trip latency. In a parallel architecture, the corresponding critical-path calculation replaces this sum; the present decomposition applies only to the specified sequential path. The network can add delay variability, retransmissions, and outages, while the compute stage can be affected by concurrent jobs, thermal management, or model-loading operations. Mean response times alone are inadequate where missed deadlines generate unsafe or defective physical actions.

For a deadline $D > 0$ (seconds), define a completion event $C_D = \{T_{\text{response}} \leq D\}$. Its probability depends on the joint distribution of component delays, including correlated congestion and equipment failures. An engineering claim of timely autonomy should specify the task, operating distribution, reliability target, instrumentation, and confidence interval. The absence of a demonstrated deadline failure in a short pilot cannot establish arbitrarily reliable deadline satisfaction. The relevant timescales might range from tightly bounded equipment feedback to much longer planning intervals, but their numerical limits must be derived from actual plant hazards and process tolerances.

## 16.4 Deterministic Communications Require End-to-End Validation

Industrial information exchange has requirements that differ from ordinary best-effort data transmission. Periodic process control may require bounded jitter and well-defined loss handling; inspection-image upload may prioritize throughput; emergency interlocks require independently justified safety architectures. NIST's guide to industrial wireless deployment emphasizes specifying factory-specific requirements and monitoring actual communications performance [19]. The published IEC/IEEE 60802:2026 time-sensitive networking profile provides configurations and interoperability guidance for industrial automation networks [57]. Its existence cannot establish that a given multi-vendor installation meets a specific end-to-end deadline without conformance, integration, and operational tests.

**Proposition 8** (Sufficient bound for a sequential computational deadline)**.** *For a fixed productive task using the sequential architecture in Eq.* (41)*, suppose each nonnegative stage delay $T_i$ satisfies $T_i \leq \bar{T}_i$ on an explicitly specified set of admissible operating conditions. If $\sum_i \bar{T}_i \leq D$, then the task-response time satisfies $T_{\text{response}} \leq D$ throughout that set.*

*Proof.* Summing the stagewise inequalities gives $T_{\text{response}} = \sum_i T_i \leq \sum_i \bar{T}_i \leq D$. No independence assumption is required. The result fails to provide an operational guarantee where one or more assumed bounds are unverified, where the execution path is misidentified, or where required safety functions extend beyond the specified response deadline. □

The proposition is an elementary sum bound, introduced only to clarify the verification obligation. Consider a *hypothetical* inspection path with tested upper bounds of 12, 20, 16, 35, and 9 milliseconds for acquisition, transmission, queueing, inference, and actuation. The sum is 92 milliseconds. A 100-millisecond deadline would be satisfied under these assumed bounds, but the eight-millisecond difference is not evidence of real-world margin unless the test conditions capture relevant overloads, retries, clock-synchronization errors, faults, and future operating changes. The result establishes neither physical safety nor the appropriateness of the deadline itself.

DRAFT

For systems with a remote planning service, a disconnected operation mode is an architectural option rather than an automatic safety solution. Whether local controllers should continue, slow, enter a validated safe state, or request human intervention depends on the specific process. Industrial redundancy should be evaluated at the task level: replicated servers sharing one power supply or communication route may have correlated failure modes. Reliability claims therefore require fault-injection tests and a documented common-cause model, not only nominal device counts.

## 16.5 Data and Model Provenance Condition Reproducible Decisions

AI-mediated production can incorporate sensor streams, maintenance histories, component specifications, model parameters, calibration certificates, human authorizations, and software images. Their relations determine whether an observed output can be reconstructed and meaningfully investigated after a quality defect or ecological incident. NIST's now-completed digital-thread project documented interoperability efforts spanning design, manufacturing, and inspection, and W3C's PROV framework provides established concepts for entities, activities, agents, and derivation [133, 196]. The proposed GR representation does not replace these approaches; it can add a research hypothesis about the consequences of recording changing feasible alternatives, access constraints, and excluded production paths where conventional records do not already preserve them.

A model-deployment record should identify its training and evaluation provenance where available, executable version, dependencies, intended validation domain, deployment date, data-retention practice, and responsible authority. Such records do not prove that the system's predictions are true. They can instead permit investigators to distinguish model drift, uncalibrated sensors, altered production inputs, unauthorized changes, and missing observations. Where commercial confidentiality, data protection, security, or other legal limits restrict disclosure, the study must document these limits and evaluate whether sufficient independent verification remains possible.

Retention also consumes physical resources. High-frequency logging can increase storage requirements, bandwidth, and computational overhead; retaining every raw event indefinitely can be infeasible. An appropriate policy differentiates safety-critical traceability, scientific reproducibility, lawful retention, and the environmental burden of the recording infrastructure. The decision is not resolved by maximizing recorded information without constraint. It requires identifying the evidence necessary for contestability and remedy, and measuring the cost and fidelity of competing recording schemes.

## 16.6 Operational Resilience Requires Cybersecurity and Repair Capacity

Computational productive systems remain susceptible to ransomware, supply-chain compromise, misconfiguration, corrupted backups, unavailable credentials, loss of remote access, and routine hardware faults. NIST SP 800-82 Revision 3 explicitly distinguishes operational technology from conventional information technology through its performance, reliability, and safety requirements [176]. Security controls that are acceptable for ordinary enterprise services may disrupt time-critical manufacturing operations if introduced without process-specific assessment. Conversely, direct exposure of production controllers to external services can create risks that must be considered within the broader physical-hazard analysis.

Evidence of resilience requires tested recovery procedures: safe isolation of affected segments, preservation and restoration of valid configurations, recovery-point and recovery-time objectives, spare equipment, operator training, and verification that restored production remains within permitted quality and safety envelopes. An outage survey gives important contextual evidence about digital infrastructure, but its samples are not direct estimates of the interruption probability for a particular factory. In Uptime Institute's 2025 analysis, 54% of respondents in its 2024 survey reported that their latest significant, serious, or severe data-centre outage cost more than USD 100,000; the sample and cost definition do not license direct extrapolation to industrial plants [30].

The reproduction problem is evident at a longer horizon. Technical independence may require recurring external security updates, specialist incident response, cryptographic trust services, and replacement

DRAFT

components. Vendor lock-in, open interfaces, contractual access, and availability of qualified personnel therefore affect effective productive options. None of these arrangements alone proves that a facility is robust; an open stack can suffer serious maintenance deficits, while a proprietary platform may provide tested service continuity. Their comparison should preserve matched hazards, budgets, operating requirements, and evidence maturity.

## 16.7 Computational Electricity Estimates Have Different System Boundaries

The electricity associated with AI production is distributed among local controllers, edge accelerators, network devices, servers, storage, data-centre infrastructure, and training or model-development systems. Global electricity statistics provide a useful measure of the expanding infrastructure but cannot directly identify energy per production decision. The IEA's 2026 assessment estimates worldwide *data-centre* consumption at 485 TWh in 2025 and projects approximately 950 TWh for 2030 in its central scenario; these figures include more than AI workloads and are not observations of industrial automation alone [68]. The 2026 Lawrence Berkeley National Laboratory update projects a central estimate of 11.8% and a scenario range of 9.5–15.3% for data centres' share of US electricity in 2030, rather than documenting a present realized share [171]. The different spatial coverage, methods, denominators, and projection horizons preclude simple addition or facility-level inference.

An appropriate accounting model distinguishes energy actually metered for a computational function from allocated overhead and embodied burdens. For workload class $w$ within site boundary $B$, a provisional accounting identity is

$$E_{B,w} = E_{\text{processor},w} + E_{\text{memory},w} + E_{\text{storage},w} + E_{\text{network},w} + E_{\text{allocated overhead},w}, \tag{42}$$

where each term is measured or estimated in kilowatt-hours for a declared period, allocation rule, and activity set. The decomposition is valid as a bookkeeping identity only if overlapping contributions are excluded: for example, network power embodied within a metered server must not be counted twice. Energy allocated by processor time, instruction count, or memory bandwidth should be reported as an estimate with sensitivity to the allocation rule. Upstream chip fabrication, equipment replacement, water use, and manufacturing supply-chain effects lie outside Eq. (42) unless an enlarged lifecycle boundary explicitly includes them. Energy and ecological damage should not be treated as interchangeable quantities.

Power-usage effectiveness (PUE) compares total data-centre facility energy to IT equipment energy for an established measurement period and boundary; it is not a measure of carbon intensity or embodied resource impact. Google's published fleet-wide trailing-twelve-month PUE was 1.09 for the fourth quarter of 2025, a company-specific observation subject to that operator's reporting conventions [49]. It does not establish the electricity efficiency of a small manufacturing edge server, the energy requirement of a particular AI inference, or the ecological viability of the company's entire infrastructure. Cross-site comparisons require consistent boundaries, equipment utilization, climate, and accounting methods.

## 16.8 Comparative Tests of Computational Productive Capacity

The proposed empirical programme compares three architectures executing the same industrial function: (i) a competent conventional local controller supplemented by ordinary site services; (ii) local or edge AI inference with validated fallback; and (iii) remote or hybrid AI services with explicitly recorded network dependencies. Each architecture should be evaluated with identical order specifications, process disturbances, test materials, inspection standards, authorization conditions, and monitoring periods wherever the comparison is feasible. Because changes in computation can alter the control policy itself, any unavoidable policy differences must be declared and examined separately.

Primary outcomes include task completion within declared deadlines; output conformity; quality of state estimates; human intervention; unplanned interruptions; tested recovery after power, network, and software faults; and the ability to restore functioning from preserved models, settings, and component

DRAFT

Table 24: Selected evidence on the physical infrastructure of production computing. Sources report different populations and should not be combined into a single estimate of industrial-AI energy.

| Evidence and reference period | Quantitative or institutional observation | Applicability and limit |
|---|---|---|
| IEA world data centres, 2025/2030 | 485 TWh estimated in 2025; 950 TWh projected in 2030 | Total data-centre electricity, not an industrial-AI subtotal [68] |
| LBNL US data centres, 2030 | Scenario range 9.5–15.3% of US electricity | Projected share, dependent on bottom-up equipment and cooling assumptions [171] |
| Google fleet, Q4 2025 | Trailing-12-month PUE 1.09 | Operator-specific facility/IT ratio, not embodied impact [49] |
| Uptime survey, respondents reporting 2024 outages | 54% reported latest significant or worse outage above USD 100,000 | Surveyed data-centre incidents, not plant-level failure probability [30] |
| IEC/IEEE 60802, June 2026 | Published time-sensitive networking profile | A standard defining profiles, not proof of deployment conformance [57] |

inventories. Resource outcomes include metered electricity, site-level water use where relevant, equipment consumption, storage/network overhead, and lifecycle burdens modeled under a transparent system boundary. The analysis should report measurement uncertainty and separate measured field outcomes from simulation or vendor specification. The comparators need not be limited to AI versus non-AI: competing industrial systems may already implement sophisticated deterministic control or optimization.

Three patterns would be substantively informative. Edge inference might improve deadline compliance while worsening local maintenance demands or underutilized-hardware energy burdens. A cloud-based architecture might reduce local hardware inventories while increasing sensitivity to communication interruptions or provider access rules. A hybrid architecture might exhibit lower operational risk if independent local safety functions remain available, while still showing concentration in long-run model-maintenance capacity. None is an inevitable consequence of architecture labels. Negative or mixed results would help specify the boundary within which a proposed GR dependency representation contributes information beyond established process maps, reliability models, input–output analysis, and standard provenance logs.

## 16.9 Implications for the Historical Organization of Production

Computational infrastructure makes the social character of ostensibly autonomous production visible. Factory decisions depend upon the accumulated work of chip designers, operators, maintenance technicians, software developers, data workers, electricity producers, and institutions governing access and repair. As production becomes more computationally adaptive, control of processing capacity, data interfaces, model updates, and interoperable technical standards may affect the range of production trajectories accessible to different actors. Such influence is contingent on substitutability, enforceable rights, redundancy, and historically specific relations of production, as investigated in Sections 10 and 11. Computing architecture therefore warrants analysis as an evolving condition of productive reproduction rather than an immaterial layer above it.

The empirical task is to determine when computational changes reduce dependencies, when they relocate them, and when they generate new concentrations or demands for specialized labour and infrastructure. The present section provides bounded technical descriptions, elementary formal conditions, and measurement proposals. It does not establish that any AI-enabled production system is self-reproducing, ecologically compatible, or governed legitimately merely because its information systems are highly automated. The next section examines the energetic and thermal infrastructures underlying both computation and the wider productive installation, where competing uses of power, cooling, and water require additional physical accounting and ecological evaluation.

DRAFT

# 17 Powering and Cooling Autonomous Production Infrastructure

The previous section identified the power demand, timing requirements, and physical replacement needs of industrial computation. This section examines the wider energetic and thermal conditions under which an increasingly autonomous production installation can operate and continue operating. Its objective is to connect electricity-system dependence, process heat, cooling, water use, backup capacity, and ecological burdens within declared physical and institutional boundaries. It considers the capability to maintain critical processes during disturbances separately from annual energy efficiency and from ecological acceptability. An AI-mediated facility may optimize its electricity schedule while relying on a constrained regional network, an unobserved upstream fuel system, or a water-intensive cooling arrangement. Equally, industrial electrification and better thermal integration may improve some of these dependencies. Neither direction is assumed in advance.

The empirical context has changed rapidly. The International Energy Agency's April 2026 assessment estimated world data-centre electricity consumption at approximately 485 TWh in 2025, with a central projection of 950 TWh in 2030; the latter is a projection for all data-centre workloads, not an observed AI-only demand [68]. Lawrence Berkeley National Laboratory's June 2026 update estimates a reference-case 649 TWh of *US* data-centre consumption in 2030 and constructs compounded-uncertainty scenarios spanning 521–843 TWh; this is a model-dependent scenario range, not a statistical confidence interval. The territorial and modeling boundaries differ from the IEA global series [171]. Neither result directly measures energy used by AI-controlled manufacturing plants. These studies nonetheless demonstrate why productive autonomy increasingly has to be assessed alongside the supply, financing, and operation of energy infrastructure.

## 17.1 Production Energy Extends Beyond the Computational Workload

A manufacturing installation may simultaneously require electricity for motion and drives, computing, air compression, ventilation, pumps, chilled-water systems, lighting, materials transport, metrology, and safety equipment, as well as fuel or electricity for process heat. Its energy requirements depend on product design, throughput, operating schedules, ambient conditions, equipment age, and rework. The electricity to operate an AI model is therefore only one term of a larger production balance. A model that improves inspection accuracy can still increase site electricity requirements if the improvement requires additional accelerators, cooling, and redundant networking; conversely, better scheduling or lower reject rates may reduce energy per conforming unit despite extra computing. Demonstrating either effect requires matched output quality and a complete allocation of added and avoided activities.

For a boundary $B$ and observation interval $[0, H]$, define the *site electricity drawn* in kWh as

$$E^{\text{site}}_{B,H} = \int_0^H \left( P_{\text{process}}(t) + P_{\text{computing}}(t) + P_{\text{cool}}(t) + P_{\text{aux}}(t) \right) \mathrm{d}t, \tag{43}$$

where $t$ is measured in hours, the $P$ terms are average electrical powers in kW at nonoverlapping meters, and on-site generation and battery flows are allocated under a separate import/export balance. This expression is an allocation identity for electricity *used by equipment*; it is not automatically the electricity *purchased from the grid*. Thermal fuel inputs, embodied energy in replacement machinery, transport, and upstream generation losses are outside its site-meter boundary. For production comparisons, the numerator must also be related to the quantity and conformity of useful output; otherwise a low-output, idling facility can appear favourable under a total-energy metric while performing poorly per accepted product.

The IEA's *Energy Efficiency 2025* review reports that industrial energy demand remains heterogeneous, with high-temperature thermal requirements dominating some energy-intensive activities and electrical motor systems important in less energy-intensive sectors [64]. Its December 2025 study of renewables for industry identifies considerable potential for electrifying low-temperature process heat and steam, while also documenting the significance of existing fuels and infrastructure [66]. Consequently,

DRAFT

no ratio derived from a computer-centre benchmark can safely be applied to a furnace, precision foundry, cold-storage facility, or robotic assembly line without process-specific measurements.

## 17.2 Electricity Supply Is a Time-Dependent Production Condition

Annual energy consumption obscures the time structure of industrial dependence. A plant with average demand of one megawatt may have much higher short-duration power peaks from equipment start-up, thermal cycles, actuation, or simultaneous processing. Electrical continuity requires adequate connection capacity, voltage and frequency quality, protective coordination, and controlled recovery after an interruption. Grid access is itself governed by planning, connection agreements, tariffs, utility procedures, and sometimes curtailment conditions. The IEA's 2026 discussion of energy and AI identifies bottlenecks in transformers, generation, grid connections, and related supply chains [68]. Those findings concern evolving electricity infrastructure and cannot establish an individual site's likelihood of interruption.

A production-facility reliability analysis should distinguish: loss of external supply; reduced firm import capacity; internally caused equipment failures; poor power quality; loss of cooling during otherwise continuous electrical service; and failed restoration after an outage. These events can have different probabilities, causes, and consequences. A semiconductor fabrication tool, a robotic manipulator, and an inventory-management server need not share a common tolerance for interruption. Some processes can pause safely, others require controlled shutdown, and still others can incur irreversible work-in-progress losses. The safety, quality, and economic consequences therefore depend on the actual production task rather than on the generic presence of a battery or redundant feeder.

On-site photovoltaic generation, fuel-based generation, batteries, thermal storage, demand response, and noncritical-load shedding can alter these dependencies. Storage is not a fuel source, and local generation does not establish islanded operability unless switching, grid-forming capability, protection, controls, and black-start procedures have been demonstrated. The US National Laboratory of the Rockies describes black-start research with inverter-based resources, including grid-forming requirements, challenging inrush loads, and testing needed for reliable restoration [134]. Its public programme description is not itself a completed performance test of any specified industrial plant, and does not establish that such a plant can island or restart safely. Backup systems themselves require inspection, replacement components, testing, fuel or charging, and trained operators. A locally resilient electrical design can therefore shift long-run dependency toward battery materials, power electronics, or maintenance institutions, rather than eliminating it.

## 17.3 Heat Balance Constrains Continuous Operation

Electrical equipment dissipates heat through resistive losses and other irreversible processes, while industrial processing also involves useful heat, latent and sensible loads, mechanical friction, and exchanges with processed materials. A cooling system must dispose of heat at a sufficient rate and maintain acceptable temperatures across equipment, materials, and workspace boundaries. Simplifying all thermal demands into a fixed percentage of electrical load conceals variability in ambient weather, process temperature, fluid properties, facility layout, and heat-transfer limitations. Air cooling, water-cooled heat exchangers, direct liquid cooling, refrigeration, and heat-pump systems have different operating domains and failure modes.

For a declared thermal zone, a lumped first-law balance may be written as

$$\frac{\mathrm{d}U_{\mathrm{th}}}{\mathrm{d}t} = \dot{Q}_{\mathrm{generated}} + \dot{Q}_{\mathrm{in}} - \dot{Q}_{\mathrm{recovered}} - \dot{Q}_{\mathrm{rejected}} - \dot{Q}_{\mathrm{out}}, \tag{44}$$

where $U_{\mathrm{th}}$ is stored thermal energy in kWh and each dotted term is a heat-transfer rate in kW, with its sign fixed by the zone boundary. "Recovered" heat is heat usefully exported to a recipient; "rejected" heat crosses the boundary without that specified use. This is a schematic control-volume representation, not a substitute for spatial temperature fields, coolant mass flow, heat-exchanger approach temperatures, or

DRAFT

transient plant models. When the critical load generates heat during an electrical outage, the continued presence of electrical power alone may fail to preserve operation if refrigeration, coolant circulation, or heat rejection is lost.

The US Department of Energy's 2024 design guide discusses air management, cooling-system design, heat recovery, and benchmarking metrics as interdependent engineering choices [187]. These are design options, not a claim that one cooling approach always minimizes environmental impacts. A dry cooler may conserve on-site water yet require greater fan power or face difficult conditions during hot weather; an evaporative system can reject heat efficiently while consuming water; a direct-to-chip liquid system can reduce some thermal resistances but still requires a downstream heat-rejection pathway. For heterogeneous manufacturing, cooling choices must also account for tool and product temperature tolerance, humidity, contamination limits, and occupational exposure.

## 17.4 Water Accounting Requires Hydrological and Spatial Boundaries

Cooling-water analysis must distinguish *withdrawal* from a source, *consumption* that reduces water available for immediate local reuse, return flow, recirculation inside the facility, and altered water quality. Cooling towers often consume water through evaporation and discharge a separate blowdown stream to manage dissolved solids; the US Department of Energy describes these mechanisms and the importance of cycles of concentration [185]. Recirculating coolant within a sealed loop does not by itself establish low total water consumption, because the external heat-rejection system and purchased electricity may have water dependencies elsewhere. Site water is also distinct from life-cycle water associated with mining, component manufacturing, and energy generation.

For computing facilities, the *water usage effectiveness* indicator (WUE) relates a protocol-defined measure of operational site-water input or use to information-technology equipment energy, under specified measurement categories and matched reporting periods [97, 185]. It is a site-operation indicator, not a cradle-to-grave water-footprint result or an indicator of local water scarcity. If the quantities refer to the same period, its units are L/kWh:

$$\mathrm{WUE}_{B,H} = \frac{W^{\mathrm{site,input}}_{B,H}\ (\mathrm{L})}{E^{\mathrm{IT}}_{B,H}\ (\mathrm{kWh})}. \tag{45}$$

Here $W^{\mathrm{site,input}}$ denotes the water-input quantity specified by the reporting protocol; it is not automatically identical to net hydrological consumption after return flows. The denominator is IT energy, not all production electricity. WUE may be appropriate for a dedicated data centre but is not a ready-made measure of water efficiency for a factory containing wet chemical processing, equipment cleaning, or product-associated water. Such a factory requires a disaggregated water balance with separately reported source, quantity, quality, discharge, reuse, and production function. In all cases water intensity is insufficient to determine local ecological significance: identical litres can have different consequences under different catchment conditions, seasons, withdrawal rules, and ecological flows.

The European Commission's 2026 report on its first EU data-centre reporting period defines WUE through total water input divided by IT equipment energy and gives an average reported 2024 WUE of 0.58 L/kWh, with Member-State reported values ranging from 0.07 to 1.28 L/kWh [35]. This is a summary of facilities in the reporting system, with reporting coverage, technology, and climate differences; it is neither a complete census of all EU data centres nor an uncertainty interval for a European population mean. In particular, it is not the lifecycle water footprint of AI manufacturing. For a purely *hypothetical* IT facility using 1,000 kWh of IT electricity and a WUE of 0.50 L/kWh over its matched measurement period, the corresponding site-water figure would be 500 L. It would not include indirect electricity-generation water, external component supply, or the freshwater-stress context. Quantitative comparison of cooling choices therefore requires reporting both on-site and broader supply-chain assumptions, without adding incompatible withdrawal and consumption quantities.

DRAFT

Table 25: Energetic, thermal, and water measurements required to distinguish productive continuity from ecological assessment. Numerical examples are not interchangeable across boundaries.

| Research object | Primary observation and unit | Principal inference limitation |
|---|---|---|
| Plant electricity | Time-stamped import, on-site generation and load (kW; kWh) | Energy sums alone omit power quality, curtailment and external generation mix |
| Supply continuity | Outages, voltage and frequency events, restoration time | Equipment redundancy does not prove critical production survives an event |
| Heat management | Heat load, temperatures, coolant flow and thermal storage | A cooling coefficient does not establish process-temperature compliance |
| Site water | Withdrawals, evaporation, blowdown, returns and quality (L or $m^3$) | Site indicators omit regional water stress and upstream water dependencies |
| Heat recovery | Delivered useful thermal energy at specified temperatures (kWh) | Gross heat availability is not equivalent to usable or displaced energy |
| Environmental burden | Location-specific energy and water supply, equipment lifetimes | Attribution, time horizons and counterfactual baselines affect comparisons |
| Institutional access | Connection rights, water permits, allocation and curtailment rules | Physical capacity does not imply authorization or equitable access |

## 17.5 Storage and Demand Flexibility Have Bounded Guarantees

To test the independence of a critical production function from uninterrupted grid service, specify a disturbance interval $[0, H]$, the essential electrical load $p_c(t)$ in kW, and guaranteed contemporaneous supply $p_f(t)$ in kW. Define the power deficit $d(t) = [p_c(t) - p_f(t)]_+$, where $[x]_+ = \max\{x, 0\}$. The deficit is defined at the declared electrical delivery interface after conversion losses already included in the guarantee. Separate the amount of usable energy deliverable from storage (kWh), its maximal output power (kW), and the operating assumptions allowing it to deliver that power. These specifications concern a bounded service guarantee, rather than complete production-system autonomy.

**Proposition 9** (Sufficient electrical backup over a bounded disturbance)**.** *Assume that, for the entire interval $[0, H]$, the critical-load profile $p_c$ and guaranteed non-storage supply $p_f$ are measurable, nonnegative and expressed at the same electrical boundary; a properly connected storage system can dispatch its output without startup delay, network restriction or further loss beyond its declared usable-energy rating; it has initial deliverable energy $E_0$ (kWh) and dispatch limit $P_{\max}$ (kW); and no other interdependent physical constraint is binding. If*

$$P_{\max} \geq \operatorname*{ess\,sup}_{t \in [0,H]} d(t), \qquad E_0 \geq \int_0^H d(t)\,\mathrm{d}t, \tag{46}$$

*then it can meet the specified critical electrical load throughout the interval by dispatching $d(t)$.*

*Proof.* Choose storage dispatch $p_b(t) = d(t)$. By the first inequality, it satisfies the assumed instantaneous output bound almost everywhere. By the second, its cumulative delivered energy never exceeds $E_0$ on any subinterval $[0, t]$ because $d$ is nonnegative. Hence $p_f(t) + p_b(t) \geq p_c(t)$ and the assumed electrical service requirement is met throughout $[0, H]$. □

This elementary result depends on demanding assumptions. Storage efficiencies, state-of-charge limits, reserve requirements, battery ageing, thermal derating, fault response, synchronisation, and pro-

DRAFT

tection must be included when defining the energy *actually deliverable* at the interface. Moreover, satisfying the electrical condition does not establish process quality, cooling continuity, water availability, or safety certification. In a hypothetical case with a constant 200-kW deficit lasting 0.5 h, the required delivered energy is 100 kWh and the discharge power capability at least 200 kW. If an initial stored-energy quantity is rated *before* a 90% discharge efficiency, at least $100/0.90 \simeq 111.1$ kWh would have to be available internally under this simplified conversion model. No real facility, battery specification, or outage frequency is inferred from the calculation.

The same modelling discipline is necessary for demand response and thermal storage. Shifting a noncritical task to a later hour can reduce peak imports, but only when production deadlines, inventory, labour schedules, ambient-temperature constraints, and future power availability permit that shift. A battery can provide short-duration continuity while its material supply and replacement remain externally dependent. A system-level assessment should report both the immediate service improvement and the altered long-term reproduction requirements.

## 17.6 Heat Recovery Has Temperature and Recipient Constraints

A heat stream has a quantity and a temperature level. Its usefulness depends on a feasible recipient, a temperature difference sufficient for transfer, piping and heat-exchanger losses, operating hours, and the possibility that recipient demand does not coincide with source availability. Recovering 1 kWh of low-grade heat does not necessarily avoid 1 kWh of purchased high-temperature process heat or primary fuel. Thermal upgrading through a heat pump may enable an additional use but introduces electricity and infrastructure requirements. The DOE design guide identifies heat recovery as a data-centre design option, while the IEA's recent industry analysis documents distinct temperature requirements and electrification possibilities across manufacturing processes [187, 66].

A carefully bounded comparison can therefore record $Q_{\text{captured}}$, $Q_{\text{delivered}}$, and $Q_{\text{displaced}}$ separately, all in kWh of heat but defined through different measurements and counterfactuals. $Q_{\text{captured}}$ is heat entering the recovery system; $Q_{\text{delivered}}$ is heat at the user boundary; $Q_{\text{displaced}}$ concerns an avoided reference service after accounting for temperature, timing, and recipient demand. The economic value of displaced fuel or heat depends on prices and institutions, and the environmental value depends on the reference supply, emissions, water effects, and infrastructure lifecycle. A physical heat balance cannot independently adjudicate whether an allocation of these benefits among an operator, district-heating user, utility, or municipality is legitimate.

## 17.7 The Production Network Determines Indirect Energy and Ecological Burdens

A production installation can reduce site electricity demand while increasing off-site embodied energy, material processing, logistics, or replacement rates. Conversely, local electrification may reduce direct fuel combustion but shift some effects to electricity generation and network investment. Attribution requires a declared functional unit, spatial boundary, temporal resolution, supplier data, and counterfactual. An annual average electricity-generation factor can differ from the marginal effect of scheduling additional consumption during a particular grid hour. The two quantities answer different questions; they should not be substituted without explanation. Section 16 made a similar boundary distinction for computation, and Sections 9 and 8 explain why the resulting dependence can emerge in upstream productive structures.

This issue is directly relevant to historical materialism. Electrification, storage, and automated thermal control alter productive forces, but access to network capacity, water, land, equipment finance, and maintenance remains historically and institutionally organized. Operators may face incentives to minimize private tariffs while ecological costs occur in a different watershed, electricity region, or supplier community. Public or cooperative ownership can change decision rights and cost allocation without automatically resolving physical constraints. Analysis should compare institutional arrangements through observable decisions, rights, and consequences rather than assuming that a change in ownership suffices for ecological compatibility.

DRAFT

Table 26: Prospective comparisons for assessing physical and institutional conditions of energy-supported productive autonomy. Outcomes may be favourable, adverse, mixed or inconclusive.

| Comparison | Matched measurements | Findings that could challenge the proposed mechanism |
|---|---|---|
| Grid-only versus tested islanding | Critical-load survival, process quality, restoration time, total backup burdens | Local storage fails under common-mode cooling or control faults |
| Flexible versus fixed schedules | Peak demand, production delays, rejects, labour intervention, grid mix | Reduced site peak shifts emissions or resource demand without net benefit |
| Dry versus evaporative rejection | Power, site water, heat-removal service, climate and water stress | Lower site water yields materially higher off-site burdens or reduced reliability |
| Recovered versus rejected heat | Delivered thermal service, recipient load, heat-pump energy, avoided reference heat | Measured heat recovery produces little useful substitution at required temperature |
| Alternative access arrangements | Connection rights, water allocation, restrictions, costs, monitoring | Measured physical conditions remain similar despite changed ownership or contracts |

At the ecological interface, indicators of water scarcity, thermal discharge effects, air emissions, habitat occupation, and production of replacement equipment need separate validation. A reduction in PUE or WUE at a particular facility is relevant evidence about declared efficiency boundaries, but the two indicators do not establish recoverability of ecological systems or compliance with collective water allocations. The ecological-viability sections that follow will develop these questions explicitly.

## 17.8 Comparative Research on Energy Infrastructure and Productive Continuity

A feasible empirical study would compare several power-and-cooling architectures on matched production tasks and time horizons: a grid-dependent installation with ordinary backup; a facility with demand response and bounded storage; a site with diversified firm supply and tested islanding; and variants with dry, evaporative, or liquid-assisted cooling. The purpose would be to determine under which conditions each architecture improves critical-process continuity, total resource burden, and reproducibility. Comparisons should identify production quality and throughput, product mix, weather, external energy contracts, metering coverage, and equipment lifecycle as potential confounders. Storage and water systems must be tracked at suitable time resolution because annual energy totals conceal peak events and seasonal scarcity.

A comparative study should also test whether additional GR-style dependency recording improves decisions relative to a competent baseline combining existing electrical asset registers, energy-management systems, water audits, process models, and life-cycle inventories. A useful outcome would involve identifying an otherwise missed dependency, forecasting a resource-constrained stoppage, or explaining a disputed allocation of costs with traceable evidence. Equivalent or worse performance after including recording burden would weaken any claim of incremental benefit. The proposed research programme does not assume that more complex representation necessarily yields better engineering, scientific, or ethical decisions.

The section's main conclusion is conditional: increasing autonomy within production processes can increase the significance of energy and cooling interfaces without fixing their eventual institutional or ecological consequences. Technically demonstrable reliability, economic feasibility, authorized resource access, and ecological compatibility remain distinct inquiries. The next section will examine the semiconductor, mineral, and equipment-supply structures through which these energetic infrastructures themselves are produced and renewed.

DRAFT

# 18 The Supply Dependencies of Advanced Production Equipment

Sections 16 and 17 identified the computing, electrical and thermal infrastructure through which productive autonomy is exercised. This section examines how such infrastructure is manufactured, provisioned and replaced. Its research objective is to identify the *qualified material and industrial dependencies* of advanced production equipment: the interdependent capabilities required to supply semiconductors, precision machinery, high-purity inputs, sensors, drives and replacement assemblies with specifications suitable for a particular production function. The historical-materialist issue is consequential. Greater automation within an industrial establishment can coexist with increasing dependence on productive capacities distributed across many other establishments and geographical jurisdictions. The structure of those dependencies influences both the technical possibility of continued production and the institutions able to control its realization.

This analysis takes neither technological autarky nor universal supply fragility as a premise. International specialization may permit lower costs, rapid diffusion, redundancy and cumulative technical improvement. It may also introduce delays or concentrated requirements where an apparently similar product cannot be substituted in an existing process. The appropriate object of research is therefore a *specified production capability*, its required quality, scale and delivery horizon, and the feasible replacement pathways available within those conditions. Commodity trade statistics, a supplier count, or a nation's aggregate fabrication capacity are useful observations, but none alone identifies the actual ability to reproduce a particular industrial system. The surveys by the OECD on semiconductor value chains and fabrication capacity offer unusually clear evidence of that distinction [141, 143].

## 18.1 Productive Equipment Requires Functionally Qualified Inputs

An automated manufacturing line consists of more than its principal robot or processing unit. Motion stages and servo drives depend on control electronics, precision bearings, calibrated encoders, sensors, power semiconductors, tool changers, hydraulic or pneumatic elements, firmware and communications interfaces. Maintaining a process at specified tolerances can require traceable calibration, qualified tooling, process gases and filtered environments. A substitute component that fits mechanically may fail electrical, thermal, metrological, cybersecurity or process-qualification requirements. Similarly, the ability to acquire additional generic computing capacity does not entail the ability to replace a validated embedded controller with a fixed response-time requirement. Relevant dependencies must consequently be typed by the *function* they support, the technical specification they satisfy, and the time and authorization required to qualify them.

For each equipment component $j$, define its requirement record as $r_j = (\sigma_j, q_j, \tau_j, \kappa_j)$, where $\sigma_j$ denotes engineering specifications, $q_j$ a declared acceptance or conformity criterion, $\tau_j$ a latest usable delivery time, and $\kappa_j$ the relevant access, licensing or authorization conditions. This is an analytical representation, not a claim that component attributes can be reduced to four naturally measurable scalars. A potential supply path is *qualified* only if its delivered component passes the required physical and institutional tests, including any transition and revalidation work. Some criteria are categorical and some depend on measured distributions rather than binary passing outcomes. Whether to impose a hard threshold, safety factor or probabilistic acceptance level is itself an engineering and institutional decision to be documented.

This perspective extends the boundary-dependent autonomy definitions of Section 7. A factory may have adequate on-site repair personnel and inventory for a finite horizon while depending on distant processes to replenish equipment that cannot be manufactured in-house. Reciprocal dependencies also arise: semiconductor manufacturing relies on its own industry's control chips, specialized measurement systems and replacement electronics. The OECD's 2025 semiconductor mapping, drawing on input–output and trade evidence, reports substantial intra-industry intermediate flows and concentrated inputs; these are observations of interconnectedness, not proof that every individual factory has a single critical supplier [141].

DRAFT

## 18.2 Semiconductor Production as a Sequence of Specialized Capabilities

A semiconductor supply network spans design tools and intellectual property, fabrication, lithography, deposition, etching, cleaning, implantation, process control, packaging, assembly, testing and logistics. Different categories of device impose substantially different requirements. High-bandwidth memory, power semiconductors, advanced logic, analogue chips and industrial microcontrollers serve different functions; their production cannot be inferred from a national wafer-start total without disaggregating process technologies and qualification requirements. The OECD's December 2025 *Chip Landscape* explicitly identifies limited substitutability between specialized fabrication plants. In particular, a facility optimized for analogue or power chips cannot be presumed able to manufacture advanced logic or memory devices merely because it is counted among wafer-fabrication plants [143].

The technical conditions of replacement include physical design compatibility, process-node requirements, wafer size, package geometry, qualification lead time, test coverage, and production yields. Some substitutions are relatively inexpensive after engineering validation, while others entail architectural redesign, software modification or complete process requalification. Changes to an industrial AI controller can affect input timing, sensor interfacing and safety evidence as well as computing throughput. Hence the relevant question is how production continues at its accepted quality, not whether another chip exists for purchase.

The scale of semiconductor manufacturing investment motivates the inquiry but must be interpreted carefully. The SEMI July 2026 midyear forecast projects worldwide original-equipment-manufacturer semiconductor-equipment sales of US$165.9 billion in 2026, including a projected US$143.9 billion wafer-fab-equipment component; both are industry forecasts rather than observed 2026 annual totals [167]. They describe spending on semiconductor capital equipment, not the number or cost of AI-directed factories and not the quantity of materials embodied in any specific robot. In historical-materialist terms, they illustrate the scale of capital investment needed to expand a particular branch of the means of production. They cannot determine the resulting ownership relations, resource burdens, or production resilience without further evidence.

## 18.3 Reproducing Precision Manufacturing Equipment

Manufacturing a specialized semiconductor tool itself requires optical systems, precision motion, vacuum components, light sources, sensing, software, metrology and many external engineering contributions. These systems are not independent commodities awaiting final assembly: product changes and supplier specifications may co-evolve through long-term collaborative development. ASML's 2025 annual report records 535 system sales, including 48 extreme-ultraviolet lithography systems, 279 deep-ultraviolet systems and 208 metrology or inspection systems, alongside approximately 5,100 suppliers, including product- and non-product-related suppliers [6]. These are *one company's* disclosed metrics, not counts for the world lithography industry or an estimate of EUV production capacity in complete machines per year. The same report discusses reliance on highly specialized single- or limited-source components and the risks of supplier disruption. A large nominal supplier network can therefore coexist with narrow, technically specialized interfaces.

A second-order reproductive question follows. Replacing a fabrication machine depends on the production, calibration and service capacities of the equipment manufacturer; those capacities depend in turn on still other specialized tools and workers. The recursive structures developed in Section 8 capture such upstream chains. An engineering assessment should distinguish components that can be repaired locally, assemblies recoverable through manufacturer refurbishment, and components that require upstream production or certified replacement. Manufacturer-supported reuse and refurbishment may reduce particular upstream requirements, but the result depends on the returned part, verification process, remaining lifetime, and whether reuse displaces new manufacture. A corporate parts-reuse indicator cannot alone establish aggregate circular material flow or ecological recovery.

Supplier access is also an institutional relation. Purchase contracts, warranty conditions, export permissions, software entitlements, source-code access and rights to calibration data may shape feasible

DRAFT

Table 27: Selected observed indicators and forecasts relevant to advanced-equipment supply. The indicators have different denominators and do not measure the autonomy or ecological burden of any individual facility.

| Source and reference period | Reported indicator | Inferential boundary |
|---|---|---|
| OECD, 2025 mapping | Globally distributed, specialized semiconductor segments and concentrated critical inputs | Value-chain mapping; no verified plant-specific replacement time |
| OECD, December 2025 | Fabrication capacity differs by chip type and process, with limited inter-fab substitutability | Nominal capacity is not qualified substitute capacity |
| SEMI, July 2026 | US$165.9 bn forecast total equipment sales for 2026 | Industry OEM-sales forecast; not actual annual sales |
| ASML, 2025 | 535 system sales and approximately 5,100 suppliers | Single-firm reporting; does not measure alternative EUV capacity |
| IEA, 2024 mineral data | Top three refiners averaged 86% for six selected energy minerals | Aggregate mineral-basket statistic; not an AI-specific material bill |
| IEA, 2025 mineral data | Top supplier averaged 72% excluding rare earths | Different statistic and mineral basket from preceding row |
| USGS, 2026 edition | 2025 output and reserves across over 90 mineral commodities | Commodity-level evidence; purity and component suitability unobserved |

repair and substitution pathways. These arrangements can affect dependence independently of the underlying physical scarcity. Section 10 distinguished the physical criticality of a component from an actor's exercisable capacity to restrict it; the present section specifies how such distinctions can be evaluated at the level of advanced production equipment rather than through aggregate market concentration.

## 18.4 Mineral Supply Becomes Industrial Supply Through Processing

Equipment reproduction draws on metals, ceramics, polymers, high-purity gases, specialty chemicals and other materials. The mineral origin of an input is only one stage in a chain that may include concentration, refining, purification, crystal growth, wafer manufacture, chemical synthesis, component fabrication and process qualification. Thus a sufficient mass of material in the ground, or even a large national mined-production share, does not establish availability of the form, purity, location or delivery schedule required for equipment repair. Conversely, recycling, process redesign, material substitution and diversified refining may change the effective dependency structure. Geological rarity and industrial supply criticality are different questions.

The IEA's *Global Critical Minerals Outlook 2025* reports that the average share of the three leading refining countries across six selected key energy minerals rose from about 82% in 2020 to 86% in 2024 [65]. The *2026 Outlook* reports that concentration continued to increase for many materials in 2025, although rare-earth refining showed a modest countertrend; excluding rare earths, the mean share of the leading refining country reached 72% in 2025 compared with 70% in 2023 [67]. These series differ in both their mineral baskets and concentration statistics. They should not be concatenated into one time series or interpreted as the fraction of minerals used by autonomous AI systems. The USGS *Mineral Commodity Summaries 2026* provides commodity-specific world output and reserve estimates for 2025, but these aggregates require additional bills of materials and process data before they can be mapped to AI-manufacturing equipment [182].

Table 27 gives the empirical scales and limitations of these sources. A future study must also distinguish intentional stockholding and legally accessible inventory from inaccessible or unqualified stock; an apparent warehouse reserve may not be available when a production fault arises.

DRAFT

## 18.5 Substitution Requires Qualification and Transition Capacity

The substitution of a supplier, component or production process can involve more than matching a catalogue description. For a specified product or operating service $p$, let $\mathcal{S}_p(H)$ denote the set of suppliers or production pathways capable of delivering a qualified input within horizon $H$. Membership depends on accepted material properties, interface compatibility, yield, quality assurance, time to qualify, institutionally available rights and logistical access. The set is therefore potentially smaller than the list of nominal suppliers. If a new supplier must undergo validation lasting longer than the remaining inventory cover, it may contribute to medium-term adaptation without ensuring immediate continuity. An alternative process may also draw on the same scarce high-purity chemical or semiconductor-tool interface, leaving an upstream common dependency unchanged.

For one simplified *serial* production chain of $m$ stages, index stage 1 as the earliest and stage $m$ as the final accepted-product stage. Let $y_i \in (0, 1]$ represent its fixed fraction of accepted output from units entering stage $i$, assume one-to-one units at each stage, and hold the stage boundaries and specification constant. For final accepted output $Q$ units over a common period, the necessary input to stage $i$ is

$$N_i(Q) = \frac{Q}{\prod_{k=i}^{m} y_k}, \tag{47}$$

when no stage has alternative routes, rework, inventory drawdown, or co-products. Units of $N_i$ and $Q$ are counts for the declared common reference product. This is elementary yield accounting, not a representation of an actual fab's nonlinear product mix. Its relevance is to show how a downstream throughput requirement becomes an upstream component and equipment demand once yields and quality restrictions are included.

**Proposition 10** (Necessary stage-capacity condition in a qualified serial chain)**.** *Under the preceding serial-chain assumptions, suppose stage $i$ can process at most $C_i$ input units during the reference period, and the available production route contains no inventory drawdown, alternative route or interstage rework. A necessary condition for $Q$ final conforming units is*

$$C_i \geq \frac{Q}{\prod_{k=i}^{m} y_k} \quad \textit{for every } i = 1, \ldots, m. \tag{48}$$

*Proof.* Stage $m$ must receive at least $Q/y_m$ inputs to yield $Q$ accepted outputs. Stage $m-1$ must supply at least $(Q/y_m)/y_{m-1}$ accepted interstage units. Repeating the same calculation upstream gives Equation (47). Each stage must process at least that many inputs, so a stage with smaller capacity cannot meet the specified final requirement. □

The proposition is a standard necessary throughput bound, *not* a claim that satisfying all inequalities is sufficient for practical production. Real industries face inventories, discrete batches, machine availability, heterogeneous product mixes, stochastic yields and correlated disruptions. A suitable expanded model would represent time-resolved capacity, inventories, process qualifications, transport, alternate routes, multi-output technologies and maintenance contingencies. Its value would lie in estimating whether replacement pathways actually exist under declared evidence, rather than in assigning an abstract score to all supply dependence.

As an explicitly hypothetical illustration, consider three serial processing stages with respective acceptance fractions $y_1 = 0.90$, $y_2 = 0.95$ and $y_3 = 0.98$. For 900 final conforming units, the first stage needs at least $900/(0.90 \times 0.95 \times 0.98) \simeq 1{,}074$ input units under the simplified assumptions. A first-stage processing capacity of 1,000 units would therefore make the target infeasible even if all later machines had ample nominal capacity. This constructed example establishes no real semiconductor yield or line-performance estimate.

## 18.6 Replacement Time Determines Short-Horizon Continuity

Material replacement has a temporal structure. For a component $j$ supporting a critical service, let $L_j$ denote qualified replenishment lead time, $S_j$ available conforming inventory, and $D_j(t)$ time-dependent

DRAFT

consumption or failure-induced replacement demand. Inventory adequacy depends on the cumulative requirement during the interval before replenishment becomes available, including uncertainty in demand, inspection failures and transportation. One useful descriptive coverage measure is $S_j/\bar{D}_j$ when an appropriate positive average demand rate is justified, but episodic failures and batch deliveries often invalidate deterministic days-of-cover interpretations. For repairable components, part-return loops, tests and refurbishment lead times must be recorded explicitly.

Resilience should consequently be evaluated by production function and recovery time, not by the location of a warehouse or headline supplier share. A production system may tolerate the delayed delivery of noncritical replacement tools while facing immediate loss of quality when a unique metrology instrument fails. Spare stock can improve short-horizon operational continuity but may shift capital, space, energy and obsolescence requirements into inventory management. Additional suppliers can improve resilience if their processes are independently capable and qualified, yet multiple legal entities may share the same fab, sub-tier contractor, transport corridor or critical raw input. Supplier diversification must therefore be assessed at the level of independently functioning pathways.

The OECD supply-chain resilience review stresses that localized production or reshoring does not by itself guarantee resilience; policy analyses require attention to trade, diversification and underlying network dependencies [142]. For advanced equipment, the practical test is whether a particular facility can recover a specified capability under relevant disruptions within its maintenance and production horizon. The counterfactual includes validated component inventories, acceptable substitution designs and authorized service capabilities, rather than assuming a local supplier is automatically independent or fast.

## 18.7 Supply Decisions Change Relations of Production

The manufacture of advanced equipment requires accumulated scientific knowledge and highly specialized skills alongside material installations. Design rights, calibration procedures, patents, export permissions, service contracts and manufacturing know-how condition who can build, operate and repair particular instruments. These historically constituted relations of access do not mechanically follow from the physical structure of a device. A component may be physically reproducible in principle but inaccessible within a time horizon because tooling, qualified personnel, source files, warranty conditions or permission are unavailable. Conversely, licensing, open standards, inter-firm collaboration, public research capacity or cooperative development may expand feasible pathways, subject to quality and financial constraints.

This provides a concrete historical-materialist problem beyond the observation that global supply chains are complex. The means of production increasingly incorporate layers of collectively developed scientific and technical knowledge, while effective control over productive continuity depends on the institutional interfaces through which that knowledge and its material embodiment can be accessed. These arrangements can support coordinated quality assurance and investment; they can also create bargaining asymmetries or restrict repair. Their consequences should be inferred from documented contracts, technical requirements, employment structures, production outcomes and alternative capacities. A change in ownership does not guarantee more suppliers or an ecologically beneficial production method. The analysis of appropriation in Section 11 therefore remains separate from the physically testable question of whether a qualified replacement exists.

The ecological dimension is similarly open. More equipment redundancy may require additional fabricated tools, raw materials and energy. Better repairability can reduce demand for new components but may involve high-purity processing and hazardous waste. Additional localized production may reduce some transport burdens while increasing facility-related water use or material losses. These effects call for lifecycle, material-flow and ecological measurements of alternatives under a matched functional output. The appropriate benchmark is not an automatically preferred ownership arrangement or geography but a documented pathway through which productive capacity is retained with specified physical and ecological consequences.

DRAFT

Table 28: Proposed tests of advanced-equipment dependency and reproduction. Each comparison requires a defined equipment function, product specification, time horizon, and system boundary.

| Research contrast | Required measurements | Possible contrary or null result |
|---|---|---|
| Nominal versus qualified substitutes | Compatibility tests, qualification records, lead times, shared upstream input maps | Alternative suppliers prove genuinely interchangeable within the target horizon |
| On-site stock versus replenishment | Failure records, validated inventory, restoration delay, part-obsolescence rate | Inventory expansion yields no meaningful continuity improvement |
| Local versus transnational sourcing | Delivered function, costs, common-mode dependencies, transport and facility burdens | Localization shifts rather than reduces vulnerability or ecological burdens |
| Repair versus replacement | Remaining service life, safety validation, materials, energy and waste per restored function | Repair performs worse over the evaluated lifecycle |
| Ownership and access alternatives | Service rights, procurement decisions, documentation access, interruption and quality evidence | Distinct institutional regimes show similar physical continuation outcomes |

### 18.8 Empirical Identification of Material Reproduction Pathways

A comparative programme could begin with three equipment classes: an industrial motion-control assembly, an AI inference server serving a manufacturing cell, and a metrology or precision-fabrication subsystem. For each class, researchers would reconstruct its bill of materials, qualified suppliers, firmware and calibration dependencies, replacement lead times, maintenance needs, and legal access conditions. The research would pair engineering tests with available trade, corporate-reporting and input–output data. It would also record what remains unobserved, particularly sub-tier supply contracts and proprietary manufacturing yields. Results should be graded by evidence type: disclosed corporate reporting, administrative customs aggregates, verified field maintenance records, engineering qualification trials or constructed model scenarios.

Table 28 expresses the paper's empirical commitments. Comparison should normalize throughput, product quality, equipment lifetime, demand profile and disturbance severity while documenting potential confounders. The hypothesis of *dependency relocation* would gain support where increased operational autonomy is accompanied by newly consequential, qualified upstream requirements; it would be weakened in a case where automation reduces or diversifies these requirements under matched conditions. The *operational–reproductive autonomy gap* would be supported only where a system performs its declared tasks while failing predeclared continuation-readiness tests. Aggregate industry expenditure or corporate supplier counts do not establish either result.

The analysis thus supplies a bridge between the immediate infrastructure of autonomous production and its material reproduction. Computing, machine control and thermal service depend on specialized upstream productive capacities that are functionally qualified, institutionally accessible and temporally available only under identifiable conditions. These conditions can be strengthened or weakened by technological and organizational choices. Section 19 will examine maintenance and repair as concrete mechanisms for preserving those capacities, while Section 20 will study recovery and replacement across equipment lifecycles. Those investigations will test more directly whether increasing autonomous production capability can be accompanied by durable, materially accountable forms of productive continuation.

## 19 Maintaining Productive Continuity through Repair

The material reproduction of advanced productive equipment, investigated in Section 18, becomes an immediate operational problem when equipment deteriorates, a component fails, software becomes unusable, or a production process ceases to meet its required specifications. This section examines the

DRAFT

capability to restore a specified productive function after such disturbances. Its research objective is to distinguish autonomous observation of machine condition from the materially and institutionally situated processes of diagnosis, safe intervention, component restoration, functional verification and renewed production. The distinction matters for the paper's central question: a factory that autonomously schedules and operates machinery may remain dependent on technicians, replacement parts, calibration services, engineering documentation and external permissions whenever productive capacity must be repaired.

Maintenance already has substantial literatures in reliability engineering, asset management, condition monitoring and production economics. The argument here neither redefines their results as novel nor assumes that AI maintenance is inherently superior. Rather, it asks what happens when repair becomes part of the *reproduction of productive conditions*: which tasks can be performed within a declared system boundary, which dependencies remain external, and which technological and institutional changes alter the possibility of returning a productive capability to service. The survey by Jardine, Lin and Banjevic establishes the familiar sequence from condition acquisition through diagnosis and prognosis to maintenance decisions [104]; ISO 17359:2018 and ISO 55001:2024 respectively locate condition monitoring and asset management within organized procedures and lifecycle responsibilities [87, 92].

## 19.1 Failure Is Defined Relative to a Productive Function

A component can remain mechanically mobile while failing its intended productive function. A robot joint that exhibits excessive backlash may still move but cannot necessarily maintain specified positioning tolerances. A pump may deliver flow with declining efficiency while remaining nominally operational. A sensor or controller may continue reporting values while its calibration, time synchronization or software integrity is insufficient for a qualified process. Conversely, some deviations can be compensated by alternate equipment, process adjustments or planned changes in output specifications. The analytical unit is thus a *required productive service*, not simply an item's binary state of being broken or working.

Let $f$ denote a defined production function, $q_f$ its acceptance conditions and $H$ an evaluation horizon. Write $S_f(t) = 1$ when the delivered service meets $q_f$ at time $t$, and $S_f(t) = 0$ otherwise. This indicator does not require a machine to possess one universal health score; the acceptance criteria can incorporate throughput, conformity, timing and safety, subject to distinct evidence. A failure event is a transition in the assessed capacity to deliver $f$, conditional on a defined task and operating regime. Wear, intermittent sensor errors, power disturbances, communication loss, contamination, cybersecurity compromise and loss of calibration may generate different restoration demands. IEC 60812:2018 provides a recognized failure-modes-and-effects methodology for systematically examining the ways items and processes may fail [63]. Such categorization does not independently establish a failure probability for a particular plant.

Failures can also propagate through interdependent equipment. Loss of a coolant pump may require suspension of a machining cell while downstream packaging remains operational for a limited period using intermediate inventories. Repair of the pump can restore the hydraulic function without immediately restoring certified product quality if thermal drift has altered the machine's calibration. A successful intervention should therefore be specified at multiple levels: physical restoration, safe restart, validated process conformity and recovery of the required production service. Confusing these levels exaggerates both the amount of autonomy achieved and the evidence available concerning reliability.

## 19.2 Condition Monitoring Does Not Complete the Repair Process

Condition-based maintenance draws on signals such as vibration, temperature, lubricant contamination, acoustic emission, current, operating cycles and quality measurements. The usefulness of a signal depends upon sensor placement, measurement uncertainty, regime changes, a defined failure mode and the consequences of missed or false alarms. ISO 13374-1 provides general guidance for software specifications concerning the processing, communication and presentation of machine-condition information [82]. Its stated scope does not independently validate the diagnostic accuracy of a particular implementation. Information transfer and analysis are necessary to many monitoring systems, yet even a correctly

DRAFT

detected fault need not determine the location or mechanism of damage. Additional testing, physical inspection or specialist interpretation may be required.

Diagnosis concerns the identification or discrimination of faults, while prognosis estimates future degradation or a distribution of remaining useful life under specified operating conditions. These are distinct inference tasks. The systematic review by Lei and colleagues maps prognostic work from data acquisition to remaining-useful-life estimation and discusses the construction of health indicators and the limitations of available datasets [110]. A forecast may be well calibrated on one fleet yet unreliable after a change in duty cycle, materials, repair history, sensor configuration or ambient conditions. Predictions of remaining useful life are conditional on the assumed operating and maintenance policy, not immutable physical deadlines.

For a condition history $y_{0:t}$, a provisional inference can be represented as $p(d_t, \theta \mid y_{0:t}, u_{0:t})$, where $d_t$ denotes latent degradation, $\theta$ model parameters and $u$ the recorded operating history. If distinct degradation states generate observations that are indistinguishable at the available resolution, the inferred state cannot be uniquely identified from those observations alone. This epistemic limit is especially consequential when repair decisions are made automatically. A high-confidence alert produced by an uncalibrated model can generate unnecessary component replacement or unsafe deferral of intervention. Independent inspection and performance verification are therefore measurable parts of repair capability, rather than signs that the system has failed to automate enough tasks.

Published benchmark data should be interpreted by generation mechanism and target population. NASA's prognostics repository includes measured bearing experiments and simulated turbofan degradation trajectories; the latter are generated by the C-MAPSS simulator, rather than field records from a deployed fleet [130]. Model performance on a simulator is useful for methodological comparisons but does not alone establish calibrated prediction, maintenance savings or production continuity in a real factory. A rigorous evaluation would additionally report alarms per operating hour, undetected failures, prediction calibration, detection-to-intervention delay, damaged components confirmed at inspection and eventual qualified restoration time.

## 19.3 Repair Requires a Sequence of Qualified Restoration Activities

Repair encompasses several potentially separable activities: detection and classification of the production loss; hazard isolation and authorization; inspection and fault localization; access to the damaged component; removal or adjustment; procurement or fabrication of parts; reassembly or software restoration; calibration; functional testing; and safe recommissioning. Some tasks may occur in parallel, and particular failures may require only a subset. The sequence must be reconstructed for the actual equipment rather than inferred from the presence of a maintenance robot or diagnostic algorithm.

A useful event record for a repair episode $e$ is

$$r_e = (f_e, d_e, \mathcal{A}_e, \ell_e, \rho_e, \kappa_e, z_e), \tag{49}$$

where $f_e$ identifies the disrupted productive function, $d_e$ the documented diagnosis, $\mathcal{A}_e$ the activities carried out, $\ell_e$ the elapsed lead times, $\rho_e$ the relevant replacement resources, $\kappa_e$ authorization and service-right constraints, and $z_e$ the evidence of verified restored service. These are typed records, not naturally commensurable numerical coordinates. The event should retain uncertainty and provenance, including whether a fault was inferred, inspected, confirmed or independently tested.

Safety is a constitutive condition of the restoration process. Maintenance personnel may face electrical, hydraulic, pneumatic, mechanical or thermal stored energy, and autonomous production controls may create unexpected movement. In relevant United States workplaces, OSHA's 29 CFR 1910.147 establishes requirements for controlling hazardous energy during covered servicing and maintenance, subject to its stated scope and exceptions [138]. The regulation is not a universal worldwide authorization rule; other jurisdictions and operations have distinct requirements. The general analytical point is that a controller's ability to command a machine cannot be presumed to include authority to bypass isolation,

DRAFT

Table 29: Evidence interfaces in restoration of a productive function. Satisfying an earlier interface does not establish success at a later interface.

| Claimed capability | Required observations | Common limitation |
|---|---|---|
| Fault detection | Time-stamped signal, defined fault, verified outcome, false-alarm rate | Alert does not establish fault cause |
| Diagnosis or prognosis | Inspection confirmation, uncertainty, performance under changed operating regimes | Prediction is conditional and may drift |
| Physical repair | Work order, component or firmware restoration, required resources, specialist actions | Replacement may leave process unqualified |
| Safe recommissioning | Isolation records, test procedures, responsible authorization, restart evidence | Technical restart differs from authorization |
| Productive continuity | Qualified output over a declared horizon, recovery duration and recurrent faults | Short run need not demonstrate sustained service |
| Reproductive autonomy | Independent repair cycles, qualified parts, skill and rights dependencies | On-site repair can remain upstream-dependent |

restart after maintenance or certify worker protection. Verification of energy isolation and responsibility for safe recommissioning can be necessary even when fault detection and physical manipulation are automated.

A comparison of repair architectures should therefore identify the interface between planned repair and authorized repair. Remote assistance, robotic inspection and digital instructions may reduce exposure or travel requirements, but some interventions demand human presence, specialist tooling, supplier-controlled diagnostics or independent sign-off. These activities are additional relations supporting productive continuation. Their presence should be measured rather than treated as residual imperfections that automatically disappear with further AI deployment.

Table 29 sets out an evidential chain through which the paper's notion of reproductive autonomy may be tested. The table's ordering is logical rather than necessarily chronological. Its central methodological contribution is to prevent success on a predictive-maintenance benchmark or in a robotic manipulation test from being reinterpreted as evidence of complete autonomous maintenance.

## 19.4 Spare Parts and Repair Lead Times Condition Recovery

The availability of replacement components links maintenance directly to the qualified-supply dependencies developed in Section 18. Repairability can be constrained by the absence of a part, the absence of a compatible substitute, a delayed supplier qualification, unavailable firmware, missing diagnostic software, or the inability to obtain design and calibration records. Such dependencies can exist even when the original machine was assembled locally. A stock of spare parts improves immediate restoration only if the parts remain conforming, accessible, correctly stored and suitable for the failed configuration.

For component class $j$, let $L_j$ denote the random time to obtain a qualified replacement, $I_j$ the conforming inventory on hand and $R_j$ the stochastic replacement requirement during a declared horizon. A simple stockout event is $\{R_j > I_j\}$ before further qualified replenishment. Estimation requires a failure process, a demand distribution and the dependence between components, not merely a deterministic average inventory ratio. Simultaneous failures arising from a common disturbance, long-tail supplier delays, or obsolescence can make independent-demand assumptions misleading. For repairable items, return and refurbishment cycles introduce additional states and durations.

Productive substitution also differs from component interchange. A controller can sometimes be replaced by reengineering the surrounding system, but this may require significant downtime, safety validation and modification of qualified process settings. Supplier-controlled firmware, access credentials and restricted service documentation may condition what is technically and legally feasible. These questions bear on the economic and institutional analysis of control over bottlenecks in Section 10; a formal right to possess a machine does not automatically convey all capabilities necessary to diagnose or

DRAFT

repair it. The paper does not presume that unrestricted service access always improves reliability: quality assurance, intellectual property, cybersecurity and specialized workmanship may support legitimate constraints, whose costs and benefits demand empirical comparison.

## 19.5 Availability Depends on Restoration as well as Failure

Reliability engineering offers an elementary way to express why reducing failures and shortening recovery are different strategies. Consider an alternating renewal model in which successive up-periods $U_1, U_2, \ldots$ and down-periods $D_1, D_2, \ldots$ constitute independent and identically distributed *cycles* with strictly positive finite mean cycle length. Assume that down periods include detection, access, physical repair, testing and all other time until the declared productive service is again available. Under the additional conditions of the renewal-reward theorem, the long-run fraction of calendar time with qualified service is

$$A_{\text{service}} = \frac{\mathbb{E}[U]}{\mathbb{E}[U] + \mathbb{E}[D]}. \tag{50}$$

The expression is a standard renewal-reward result, not a newly developed theorem. The commonly used mean-time-between-failures and mean-time-to-repair notation is only comparable when their operational definitions match these up- and down-periods; administrative downtime exclusions and censored observations can otherwise distort an availability claim.

**Proposition 11** (Conditional benefit of reduced restoration time)**.** *Suppose two policies operate on comparable alternating-renewal systems with the same finite positive mean up-period $\mu_U$ and finite mean down-periods $0 \le \mu_{D,1} \le \mu_{D,0}$. Under the renewal-reward assumptions, policy 1 has long-run service availability at least that of policy 0. Equality holds when the mean down-periods coincide.*

*Proof.* By Equation (50), availability under policy $k$ is $A_k = \mu_U/(\mu_U + \mu_{D,k})$. For fixed $\mu_U > 0$, the function $x \mapsto \mu_U/(\mu_U + x)$ is nonincreasing for $x \ge 0$ and strictly decreasing for $x > 0$. The stated ordering and equality condition follow. □

The proposition is conditional and narrow. Maintenance policy may also change failure intensity, planned downtime, product quality, workforce exposure and the distribution of long outages. A lower mean downtime need not improve availability if it is accompanied by sufficiently shorter up-periods. The stationary fraction does not establish the probability of surviving a critical uninterrupted production interval, nor does it verify safety or ecological acceptability. For a finite horizon $H$, the empirical study should separately measure the distribution of qualified output, recovery time, stockouts and out-of-service events that exceed declared thresholds.

As an explicitly hypothetical accounting example, a system with mean up-period 400 hours and mean down-period 20 hours has $A_{\text{service}} = 400/420 \simeq 0.9524$ under this model. If a verified intervention lowers mean downtime to 12 hours while preserving the up-period distribution and all other model conditions, the fraction becomes $400/412 \simeq 0.9709$. This arithmetic does not predict an improvement for an existing installation, and the hypothetical numbers are not field observations.

## 19.6 Autonomous Repair Remains a Distributed Social Capacity

An industrial system can diagnose faults with software while depending on maintenance engineers to evaluate risks, access enclosed equipment, restore precision alignment and authorize a restart. The relevant labour includes work by operators, mechanics, electricians, production engineers, metrologists, cleaners, software technicians and supplier-service staff. Much of it is performed outside the immediate automated work cell, and therefore may disappear from a narrow calculation of direct labour-hours per manufactured unit. Section 12 already showed why displacement from task execution does not determine total employment or the reproduction of human capabilities. Maintenance offers a concrete mechanism: reducing operator intervention can intensify requirements for diagnosis, specialized training, remote support, hazardous-energy control and response to rare but consequential faults.

DRAFT

For research purposes, maintenance autonomy should be represented as a conditional ability to complete a specified restoration task without particular categories of external intervention, given declared equipment, materials, information, skills and permissions. The actor performing a task may be an on-site robot, a remote human, a contracted service organization or a specialized supplier. Their actions should not be collapsed into an undifferentiated notion of "the system" when evaluating autonomy. A remote expert who provides a decisive repair instruction is part of the reproduction relation even if no one enters the production cell. Conversely, structured human oversight may improve safety and learning and should not automatically be counted as a failure of responsible automation.

The production relations governing maintenance influence this configuration. Contracts can allocate access to machine diagnostics, trained personnel, replacement components and service documentation. Public procurement, cooperative workshops, manufacturer support agreements and open technical standards can alter those access routes, but no arrangement guarantees short recovery times or proper qualification without supporting capabilities. ISO 55001:2024 establishes general asset-management-system requirements focused on lifecycle performance, risk and expenditure; it is not a certification of repair autonomy for a specific AI-operated factory [92]. An empirical comparative study should record who can authorize service, how knowledge is transferred, whether workers can contest unsafe restart decisions and which party assumes long-term maintenance obligations.

## 19.7 Ecological Consequences of Repair Choices Require Lifecycle Evidence

Repair can extend equipment service life and reduce replacement demand, but the direction of its ecological effect depends on the functional unit and the comparison. A repair operation may require energy, cleaning agents, replacement bearings, transport, process gases, wastewater management and hazardous-material handling. Continued use of inefficient equipment can also increase operational resource demand compared with a qualified new system. Conversely, premature equipment replacement may impose embodied manufacturing burdens that exceed any savings during the evaluated horizon. These possibilities cannot be resolved by assigning repair an intrinsically beneficial or harmful ecological status.

A suitable comparison defines the service delivered, required precision, operating horizon and lifecycle boundary, then records cumulative production of new parts, refurbishment energy, transport, waste treatment, operating consumption and relevant ecological effects. Common-mode outages and redundant equipment may alter the inventory required to achieve the same service reliability. The analysis should also consider whether repair shifts burdens to other establishments, including mining, precision manufacturing, solvent treatment and e-waste processing. The material-circulation analysis developed in the next section will extend this reasoning to end-of-life and recovery pathways. No ecological conclusion follows from the availability proposition above: uptime is a technical service measure, and cannot substitute for material-flow or ecological evidence.

## 19.8 Comparative Research Designs for Restored Production

The research programme should compare maintenance regimes against competent existing practice instead of assuming an advantage for AI prediction. Candidate regimes include fixed-interval preventive maintenance, condition-based inspection, validated predictive maintenance and AI-assisted fault diagnosis with different repair and authorization arrangements. Randomized deployment may be infeasible for essential or hazardous production; staggered introduction, matched facilities, matched equipment cohorts and carefully specified interrupted time series are possible alternatives when their identification assumptions can be justified. Models should allow differences in operating loads, age, supplier support, workforce experience, fault severity and production demands. Events occurring after a maintenance-policy change cannot automatically be attributed to that policy.

A suitable pilot could reconstruct repair episodes from two classes of equipment: a precision robotic assembly cell and a production-support computing or cooling unit. For each episode, investigators would record timestamped detection, diagnosis, work authorization, parts and skill requirements, isolation, physical intervention, verified restart and recurrence. The outcome would be the recovery of a

DRAFT

Table 30: Comparative investigation of maintenance and repair. Outcomes refer to defined production functions and observation horizons, not machine counts alone.

| Research comparison | Measurements and design conditions | Contrary, neutral or mixed result |
|---|---|---|
| Forecasting versus preventive inspection | Verified failures, prediction calibration, alarms, workloads, matched operating conditions | Better forecast score does not shorten recovery |
| Robot-assisted versus conventional repair | Time from isolation to tested recovery; part and specialist inputs; injury exposure | Robot assistance adds integration or validation delays |
| Open versus vendor-mediated service | Qualified documentation access, firmware controls, response times and conformity | Vendor support restores more reliably in the relevant setting |
| Spare stock versus replenishment | Qualified stock, obsolescence, demand correlation, shortages and recovery tails | Inventory growth adds burden without continuity benefit |
| Repair versus equipment replacement | Functional-equivalent lifecycle energy, material inputs, waste, cost and quality | Replacement has lower total impact or better continuity |

declared service under its relevant constraints. The study would distinguish observed failure records from artificial fault injection and simulator trajectories. It should also document missing records and selection effects: costly repairs may be more thoroughly recorded than minor adjustments, and successful interventions may be more likely to appear in vendor case reports.

Table 30 specifies comparisons that could strengthen or weaken the paper's hypotheses. Evidence of increasing diagnostic autonomy accompanied by stable dependence on external repair resources would support a particular operational–reproductive divergence. Conversely, independently verified repairs with shorter recovery time and no additional unrecorded support could narrow that gap under the specified boundary. Neither outcome establishes general self-sufficiency or ecological compatibility. The empirical interpretation must additionally examine demand effects, material flows, labour organization and the institutional conditions determining who can restore and control productive capability.

Repair thus provides a concrete test of the historical reproduction problem. The effective continuity of an AI-mediated production system depends upon a succession of material and epistemic tasks: recognizing degraded function, obtaining reliable evidence, safely intervening, acquiring qualified resources and proving the restored capability. Automation may transform these relations, but its consequences cannot be established by the mere presence of predictive software or industrial robots. Section 20 will examine how repair interacts with remanufacturing, material recovery and replacement across equipment lifecycles, thereby extending the analysis from restoration episodes to the circulation of productive materials.

# 20 Recovering Productive Materials across Equipment Lifecycles

The preceding section examined the restoration of productive service after equipment failure. This section extends the inquiry to the successive material pathways through which production equipment is retained, reused, refurbished, remanufactured, disassembled, recycled or discarded. Its purpose is to determine what material recovery contributes to the *reproduction of productive capacity*, and to identify the technical, institutional and ecological conditions under which recovered components or substances can actually substitute for newly supplied ones. In an AI-mediated manufacturing system, an automated decision to replace a controller or machine is consequential well beyond the local production cell: the decision affects equipment lifetimes, the quality and location of recoverable stocks, the demand for qualified suppliers, and the labour and infrastructures through which recovered materials become usable once again.

Circular-economy accounting and industrial ecology already provide substantial methods for investigating these questions. The proposed generative-relational (GR) contribution is therefore narrowly defined: to retain the distinctions among material conservation, *qualified functional replacement*, rights

DRAFT

Table 31: Distinct claims associated with equipment-lifecycle recovery. The evidential standard changes with the output whose recovery is claimed.

| Pathway and output | Evidence needed for a recovery claim | What the claim alone cannot establish |
|---|---|---|
| Reuse of an assembly | Serial identity, operating history, functional and safety qualification | Equivalent remaining life or lifecycle advantage |
| Repair or refurbishment | Fault and restoration records, acceptance testing, retained parts | Reproduction of the whole equipment platform |
| Remanufacture of equipment | Disassembly and processing records, conformance testing, parts traceability | Independence from external tooling or supply |
| Recycling to secondary feedstock | Collected input, separate process yields, purity and usable secondary output | Interchangeability with a qualified industrial input |
| Energy recovery and disposal | Mass destinations, emissions and treatment conditions | Retention of material or machine functionality |

over assets and repair information, and the changing productive possibilities made feasible by each recovery pathway. ISO 59020:2024 specifies a framework for measuring circularity at defined product, organizational and economic-system boundaries; ISO 14040:2006 and ISO 14044:2006 establish familiar life-cycle-assessment principles and requirements [93, 84, 83]. Circularity-indicator results do not, solely by virtue of their form, measure net ecological improvement: associated lifecycle burdens and site-specific effects require separate assessment. These approaches remain necessary baselines. A detailed dependency record is useful only if it resolves an identification or comparison problem that the competent application of those methods would otherwise leave unresolved.

## 20.1 Recovery Pathways Preserve Different Productive Qualities

The term *recovery* covers outcomes with different material and functional meanings. Reuse transfers an item or assembly to another service with little or no transformation. Repair restores a particular defective function; refurbishment may renew several features without establishing equivalent performance to a newly manufactured system. Remanufacturing generally entails a more comprehensive process of disassembly, restoration, reassembly and qualification to a declared product or process specification. Recycling processes discarded material into secondary raw materials, while energy recovery converts suitable residual fractions into usable energy rather than preserving their material functionality. Disposal retains neither a useful product function nor necessarily recoverable feedstock. These definitions require declared acceptance standards and must not be interpreted as mutually exclusive stages of a universal hierarchy: a recovered machine can contain remanufactured assemblies, replaced electronics and recycled scrap simultaneously.

The relevant measure is the *productive service* that a recovered object can deliver. A motor may contain recoverable copper and steel while its bearings, insulation or controller are unsuitable for direct reuse. Recovering nearly all the unit's mass does not establish that the original motor function has been reproduced. Conversely, a functionally qualified replacement assembly may preserve a high-value manufacturing capability while recovering only part of the original unit's mass. Whether one strategy is ecologically preferable also depends on energy, chemicals, emissions, future operation and the opportunity costs of alternative service pathways.

Table 31 specifies the evidential interfaces through which statements about recovery should pass. This prevents a reported collection rate, a high elemental recovery fraction or a successful robotic disassembly demonstration from being converted into an unsupported assertion of complete productive reproduction. Research on the remanufacturing of heavy stamping machinery, for example, investigates the reuse of structural elements through case-specific cost and lead-time assessment. It provides a grounded illustration of retained equipment capacity, while its findings cannot be transferred automatically to semiconductor equipment or laboratory instruments with different tolerance and contamination requirements [23].

DRAFT

## 20.2 Collection and Disassembly Determine the Recoverable Stock

Equipment cannot be recovered merely because its materials remain physically present. Identification, collection, segregation, safe discharge, dismantling, transport, contamination control and ownership transfer are often necessary before a component or material enters a feasible recovery process. Complex products can combine metals, engineered plastics, batteries, coolants, adhesives, coatings, glass, lubricants and substances of concern. A processing method suitable for one subassembly may be inappropriate for another; a material stream can also lose functional quality when components are mixed before separation. Recovery rates therefore depend on design for disassembly, available tooling, component information, collection logistics and the treatment actually applied.

The *Global E-waste Monitor 2024*, published by ITU and UNITAR, estimates that the world generated 62 million tonnes of electronic waste in 2022 and documented 13.8 million tonnes, or 22.3%, as formally collected and recycled [101]. The 22.3% indicator has the mass of generated electronic waste as its denominator. It is not the metallurgical recovery yield of an individual facility, the percentage of robots recovered, or proof that discarded semiconductors can be regenerated as qualified new devices. Electronics from productive equipment constitute only part of the much broader reported waste population. The monitor also reports substantial differences between equipment classes and regions, indicating that a global average is inadequate for reconstructing a particular manufacturing network.

Recordkeeping has physical consequences. An equipment passport containing composition, manufacturing revisions, hazardous-substance disclosures, repair history, firmware compatibility and disassembly instructions may reduce uncertainty about the feasible recovery path. Such records require provenance and updating: a bill of materials prepared at manufacture can become inaccurate after undocumented repair or parts substitution. EU Regulation 2024/1781 establishes an ecodesign framework with provisions for product information and digital product passports. It does not, by itself, impose a general digital-passport obligation on every product: the covered groups, required information and compliance dates depend on applicable product-specific delegated acts or other relevant legislation [36]. It is therefore an evolving institutional mechanism, not evidence that all industrial robots or computing components already possess standardized, complete passports.

The ownership and accessibility of this information also matter. Equipment leased under a service contract, purchased outright, or maintained through a vendor may have different pathways for authorized dismantling, firmware transfer, spare-part access and resale. These rights can affect the recoverability of productive capacities even when the physical equipment is identical. They link lifecycle management to the analysis of access leverage in Section 10, but do not imply that all restricted-access arrangements prevent recovery; quality assurance and hazardous-material controls can also justify specialized handling.

## 20.3 Qualified Reuse and Remanufacturing Depend on More than Material Retention

Maintaining an assembly's functional identity can preserve fabrication effort, accumulated material transformation and calibrated productive capabilities that are lost when the assembly is reduced to constituent feedstocks. A refurbished precision stage may require replacement guides, metrology and new controllers while retaining its structure; a remanufactured drive may recover castings and machined components but need new seals and bearings. Each decision depends on wear history, tolerance requirements, testing coverage, compatibility with current interfaces and foreseeable operation. Even nominally identical used components can have different remaining-life distributions and different failure correlations after a common environmental exposure.

For a recovered unit $i$ and target function $f$, let $q_{i,f}$ denote the result of a declared qualification procedure, incorporating the measurements and categorical approvals necessary for that function. Functional substitution is admissible only when $q_{i,f}$ satisfies the specified criteria, and when the required installation, validation and authorization are available within the relevant horizon. A vendor's declaration of reusable mass, or a successful bench test of an individual part, does not determine $q_{i,f}$ for an integrated production cell. This recalls the distinction between component-level capability and qualified productive service introduced in Sections 18 and 19.

DRAFT

Certain technologies pose particularly demanding challenges. Batteries have chemistry-dependent ageing, residual energy and safety requirements; semiconductor components may depend on encapsulation, traceability, proprietary interfaces and non-recoverable microscopic structures. Harper and colleagues survey lithium-ion battery recycling and discuss collection, dismantling, processing routes, economic constraints and the preservation or loss of material value [53]. The paper is an important technical reference rather than evidence that all battery chemistries, scales or contemporary markets achieve a uniform recovery yield. Remanufacturing must also account for new components consumed during restoration, the disposal of rejected subassemblies, and the energy used to establish qualified performance.

A relevant comparison evaluates at least three alternatives for the *same declared productive service*: (i) retain and maintain an existing machine, (ii) remanufacture or refurbish it, and (iii) replace it with a new machine. The study should normalize service quality, throughput, expected operating horizon and risk tolerance. Replacement may outperform refurbishment if older equipment has sharply greater operational energy consumption or cannot meet necessary safety standards. Conversely, early replacement may impose substantial embodied production burdens and reduce access to affordable productive capacity. Neither outcome should be prescribed before lifecycle evidence is collected.

## 20.4 Recycling Has Composition-Specific Yields and Quality Constraints

The physical composition of recovered industrial equipment determines feasible separation and processing routes. Mechanical separation can generate steel, aluminium, copper-bearing material and polymer fractions, often with varying purity. Hydrometallurgical, pyrometallurgical and other chemical processes can recover selected constituents from electronic scrap or batteries but require reagents, energy, temperature control, suitable feedstock and appropriate emissions management. A high recovery rate for one element need not imply a comparable recovery rate for another element in the same device, and successive operations may create cumulative losses. Recovery of a metal in an intermediate concentrate must also be distinguished from delivering the purity and specification needed by a new manufacturer.

Critical-mineral literature suggests both the importance and the limitations of secondary supply. The IEA's *Global Critical Minerals Outlook 2026* assesses recycling as a potentially growing contribution to supplies of selected energy minerals and documents significant processing concentration [67]. Its scenarios and mineral-specific definitions should not be converted into a single present-day material-recovery coefficient for autonomous factories. An inventory must identify the actual element or alloy, recovery chain, geographic system boundary, secondary-product quality and time horizon. Closed-loop uses within the same product category and open-loop uses in different products also have different implications for the reproduction of advanced equipment.

In particular, a robot's recovered material mass can be substantial while the supply of qualified sensors, actuators, chips and precision bearings remains unchanged. This creates a possible *quality bottleneck*: material circularity increases without a corresponding increase in the feasible replacement of the highest-specification productive components. A useful empirical outcome is therefore the proportion of *functionally qualified replacement requirements* met through verified reuse, remanufacturing or secondary-material supply, reported beside recovered mass by element and associated environmental impacts. This is not a universal circularity score, and its numerator requires actual qualification evidence.

## 20.5 Material Balance and a Conditional Replenishment Bound

A conservation-based model makes the accounting issue explicit. Fix one *homogeneous, usable material grade* $g$, a production-system boundary $B$ and a discrete observation interval $t$. Let $M_t$ (kg) be usable inventory at the interval's beginning, $V_t$ (kg) virgin or other qualified external primary supply, $R_t$ (kg) secondary feedstock that *passes the same usable-grade specification*, $Q_t$ (kg) material incorporated into newly produced output, and $W_t$ (kg) other irrecoverable process losses or removals. Under a boundary that records transfers without double counting, conservation gives

$$M_{t+1} = M_t + V_t + R_t - Q_t - W_t. \tag{51}$$

DRAFT

The term "virgin" is used here as a category of new primary external supply; external deliveries of secondary feedstock belong to $R_t$ and should be labeled by their actual origin. Input categories must be separately recorded to avoid treating imported recycled metal as locally recovered material. The equality is a physical balance for one grade rather than an equation relating dollars, ecological quality or productive capability.

Let $C_t$ (kg) denote feedstock physically eligible and collected for processing into grade $g$, and let the *effective qualified yield* satisfy $0 \leq \eta_t \leq 1$ and $R_t \leq \eta_t C_t$. This coefficient includes only the fractions that ultimately meet the declared usable-grade specification; an intermediate refining yield or recovery of a chemically related substance does not suffice. When recovered material crosses periods, $R_t$ and $C_t$ must be aligned through an explicit processing inventory or cohort tracking. For the following one-period result, assume that any associated processing lag has already been accounted for.

**Proposition 12** (Necessary new supply under a bounded material stock)**.** *For the balance in Equation* (51)*, assume* $M_{t+1} \geq M_t$*,* $W_t \geq 0$*,* $Q_t \geq 0$*, and* $0 \leq R_t \leq \eta_t C_t$ *with* $0 \leq \eta_t \leq 1$*. Then the required primary input obeys*

$$V_t \geq \max\{0,\, Q_t + W_t - \eta_t C_t\}. \tag{52}$$

*In particular, if* $C_t = Q_t > 0$*,* $W_t = 0$ *and* $\eta_t < 1$*, then* $V_t \geq (1 - \eta_t) Q_t > 0$*.*

*Proof.* Rearranging Equation (51) gives $V_t = Q_t + W_t + (M_{t+1} - M_t) - R_t$. The stock condition and yield bound imply $V_t \geq Q_t + W_t - \eta_t C_t$. Since nonnegative primary input is assumed, the maximum with zero follows. The special case is immediate. □

The proposition is an elementary accounting inequality, not a novel thermodynamic theorem or an empirical estimate of worldwide mineral depletion. Its scope is deliberately narrow: a single grade, a declared interval, no unrecorded secondary stocks or grade conversion, and unchanged or expanding usable inventory. If an operator deliberately draws down inventory, imports higher-grade material, redesigns a product for lower material demand, or changes the boundary, the premises and interpretation change. Genuine recycling can lower primary demand materially even while some replenishment remains necessary.

For illustration only, suppose a hypothetical equipment programme needs $Q_t = 100$ kg of qualified input and collects $C_t = 80$ kg of compatible end-of-life feedstock. If the process delivers $\eta_t = 0.75$ qualified yield, at most 60 kg becomes usable secondary supply. With $W_t = 0$ and no inventory drawdown, at least 40 kg of primary supply is required. This calculation is not a measured robot-recycling result. If the newly produced equipment has changing material composition or if the collected stock differs by grade, a vector-valued inventory with explicit transformation matrices is required; component counts cannot be treated as kilograms.

## 20.6 Ecological Assessment Requires Service-Equivalent Counterfactuals

Circularity and ecological compatibility are related but distinct assessment objects. The use of secondary materials may avoid primary extraction, but recycling also creates energy demands, chemical exposures, transport flows, process emissions and residues. Reuse can prolong access to productive capabilities while retaining inefficient equipment. Remanufacturing may avoid some embodied processing while requiring additional spare parts and validation. A comparison that considers only material recovery per tonne can therefore miss increased operating consumption or adverse impacts at another site. ISO 14040 and ISO 14044 place goal, scope, inventory, impacts and interpretation in a coherent comparative life-cycle-assessment process [84, 83]. The research should make allocation conventions, avoided-burden assumptions and displaced production explicit and examine sensitivity to those choices.

There are also strong empirical reasons to avoid equating global collection and recycling indicators with ecological outcomes. Eurostat reports that recycled material accounted for 12.2% of material use in the European Union in 2024. Its circular material-use indicator measures the contribution of materials returned to the economy relative to total material use, rather than the recovery rate of a particular industrial component [37]. Separately, the ITU–UNITAR global electronic-waste indicator measures documented

DRAFT

formal collection and recycling against waste generation in 2022 [101]. These denominators and system boundaries differ. Both contextualize material dependence, while neither identifies the lifecycle superiority of a recovered actuator or predicts the replacement of primary input in a specified autonomous factory.

A rigorous comparative protocol could use a functional unit such as a specified quantity of conforming machined products produced over five years at a declared safety and availability level. Investigators would then record the manufacture or remanufacture of equipment, repair parts, electricity, cooling, consumables, transport, waste treatment and the environmental conditions affected across suppliers. Potential rebound effects matter: reducing the cost of replacement may increase production volume, and longer equipment service may increase cumulative throughput. A full ecological assessment further considers location-specific water availability, hazardous releases, land-use change and ecological recovery. It cannot rely solely on carbon dioxide equivalents or aggregate recovered mass.

## 20.7 Material Circulation Changes Institutional Control and Productive Possibilities

The ability to recover materials is governed by productive relations as well as separation technology. A facility may own its machines but rely on a specialist service provider to access firmware, dismantle critical modules, obtain replacement chips or verify safety. A recycling organization may possess metallurgical expertise but depend upon collection arrangements, permits and viable markets for secondary feedstock. Product passports and procurement specifications may make recovery easier to organize, yet their benefits depend on data accuracy and practical access. Different investment and ownership arrangements can also affect whether long-lived machinery, easily disassembled products or short-cycle replacement strategies are selected.

These distinctions help avoid conflating *ownership of waste* with *control over reproduction of productive capacity*. A firm can extract substantial monetary value from discarded equipment while recovering few functionally useful components; conversely, a public or cooperative repair network can retain productive capabilities with relatively modest monetized returns, provided it has qualified personnel, infrastructure and lawful access. Neither institutional category guarantees ecological improvement. The paper's historical-materialist contribution is to investigate how design standards, service rights, procurement contracts and financial incentives shape which materially possible circulation pathways are realized. The EU's 2024 ecodesign framework provides one example of institutions seeking to influence durability, repairability and remanufacturing potential, subject to specific implementing rules [36]. Consumer-oriented right-to-repair provisions should not be assumed to apply in the same way to industrial equipment [34].

A typed lifecycle record can connect the relational model of Section 6 with physical material-flow accounts. For equipment identifier $i$, record the component genealogy, material grade and mass, service and refurbishment episodes, measured recovery outputs, qualified downstream applications, transport and energy flows, and the institution authorized to act at each step. Provenance alone does not guarantee accuracy: the record should distinguish measured flows from model estimates, undocumented losses and claims that have never been tested. Competent conventional databases, life-cycle inventory systems and product passports should form the baseline against which any additional GR trajectory representation is evaluated.

## 20.8 Comparative Studies of Lifecycle Reproduction

Table 32 proposes comparisons requiring independent validation rather than a presumption that automation or circularity is beneficial. A feasible pilot would follow successive cohorts of industrial control cabinets, electric drives and precision robotic assemblies from installation through maintenance to retirement. Each item would retain a verified bill of materials, changes during service, component removals, collected fractions, recovery processes, actual qualified second uses and destination of unrecovered residues. Outcome measurement would include the same productive-service requirements before and

DRAFT

Table 32: Proposed tests of material recovery in autonomous production. Each outcome concerns a specified service, equipment cohort and system boundary.

| Comparison | Necessary measurements and controls | Result that could weaken the proposed benefit |
|---|---|---|
| Reuse or remanufacture versus replacement | Qualified output, remaining life, energy, repair inputs and failure distributions | New equipment has lower impact at equivalent service |
| Recycling versus primary feedstock | Collection mass, grade-specific process yields, losses, purity and treatment burdens | Secondary grade fails qualification or creates higher impacts |
| AI-guided versus conventional recovery | Matched feedstock, sorting accuracy, recovery yield, labour, energy and costs | Model improvement has no qualified recovery gain |
| Open versus restricted recovery information | Documentation access, disassembly quality, authorization and safety validation | Access expansion does not improve qualified throughput |
| Local versus distributed treatment | Logistics, processing emissions, capacity, recovery quality, disposal and ecological exposure | Centralized processing has lower measured total burden |

after recovery, with uncertainty in service life, collection completeness, processing yields and upstream environmental impacts.

The design should distinguish genuine reductions in primary demand from accounting shifts. A recovered material sold to another sector is not necessarily a direct replacement for the original equipment's high-specification input. Avoided production is a counterfactual claim requiring evidence about what purchasers would otherwise obtain. Functional substitution should be tested through engineering acceptance, while environmental superiority requires a matched lifecycle assessment and context-relevant ecological measurements. Missing informal treatment, undisclosed export destinations and proprietary dismantling data can systematically bias observed performance. Sensitivity analyses must therefore consider unobserved waste destinations, quality rejection, demand growth and equipment obsolescence.

The outcomes may be mixed. A programme might achieve strong mass recovery while remaining unable to reproduce specialized manufacturing capability; alternatively, remanufacturing may substantially reduce the demand for new precision assemblies while still requiring new electronics and energy-intensive processing. A third outcome could show that qualified recovery improves both material and ecological measures over the declared horizon. The present manuscript does not resolve these possibilities empirically. It establishes a research framework for documenting them and for relating material circulation to the historically structured reproduction of productive forces.

The engineering survey therefore ends with a conditional conclusion: the continued use of productive equipment depends upon material stocks, qualified transformations, accumulated production knowledge, labour, access rights and ecological conditions, while circularity describes only selected aspects of those relations. Section 21 will examine the ecological conditions of productive continuation more directly. It will ask which biophysical processes and generative capacities must remain viable when production systems expand, recover materials or reorganize their material dependencies.

# 21 The Ecological Conditions of Productive Continuation

The engineering reviews in Sections 13–20 examined the machinery, computing, electricity, specialized supply chains, maintenance, and material-recovery processes through which an AI-mediated production system may continue operating. Their results leave a further question unresolved: which ecological processes must continue, and which ecological transformations remain admissible, while the technical conditions of production are repeatedly renewed? This section investigates that question through historical materialism, contemporary ecological science, and a provisional generative-relational (GR) account of heterogeneous productive and ecological relations. Its role is to establish the ecological object of assessment before subsequent sections examine recoverability, life-cycle burdens, robust viability, and

DRAFT

cumulative impacts. Ecological continuation is considered on spatially and temporally declared terms; it is neither inferred from operational uptime nor reduced to a monetary valuation of ecosystem services.

## 21.1 Socioecological Metabolism in the Reproduction of Production

Material production reorganizes exchanges of matter and energy among working populations, technical artifacts, soils, waters, living organisms, and atmospheric and geological processes. In historical-materialist terms, the reproduction of productive forces occurs within historically constituted relations of appropriation, labour, knowledge, infrastructure, and access to material conditions. Marx's treatment of the labour process and the conditions of agricultural production offers resources for examining these exchanges, while subsequent scholarship on the metabolic rift reconstructs how economic development can interrupt the renewal of soil fertility and other environmental conditions [123, 40, 41, 162]. These historical arguments supply an interpretive foundation rather than a general ecological law. The mechanisms connecting contemporary semiconductor manufacture, automated assembly, or data-centre cooling to particular ecosystems must be established through suitable physical and biological evidence.

The reproduction of humanly organized production and the continuation of affected ecological processes need different evidence and criteria. In this manuscript, *ecological compatibility* denotes a defensible, bounded assessment of productive interventions against specified ecological endpoints, uncertainty models, and justified admissibility criteria; it does not name an intrinsic property of an AI controller or a universal status granted by efficiency alone. A firm may sustain machinery by importing water, refined minerals, and replacement components while imposing cumulative burdens on aquifers, habitats, or receiving waters. Conversely, ecological functions may persist when a facility becomes inoperable because equipment fails, trained workers depart, or an authorization expires. These possibilities establish a logical separation between the continuation of productive capacity and the continuation of ecological processes. A decision to treat a particular ecological state as an inviolable constraint requires normative and institutional justification; measurement of that state and estimation of its response to intervention require ecological science. Keeping these tasks distinct prevents the objective of production from silently defining what counts as ecological continuation.

The metabolic perspective should also avoid an exclusively instrumental account of nonhuman life. Habitats and organisms may be affected by production even where they supply no identifiable commercial input, regulating service, or foreseeable benefit to an enterprise. Ethical consideration of those lives and the institutional authority to make decisions about their habitats require further argument. The present section accordingly distinguishes three claims: ecological processes may materially condition human production; productive interventions may affect ecological beings and relations independently of those instrumental conditions; and decisions about acceptable change require situated scientific evidence and normative and legal deliberation. None of the three follows automatically from the others.

## 21.2 Ecological Generativity and Conditions of Continuation

The term *ecological generativity* is used provisionally for the evolving capacity of organisms and ecological assemblages to sustain, renew, and transform living relations, subject to historically situated biophysical constraints. It is an organizing concept, not an empirically standardized scalar measure or a synonym for resilience, biodiversity, or biomass. It draws upon established ecological accounts of reproduction, trophic interactions, nutrient cycling, succession, disturbance, and resilience rather than offering an alternative biological mechanism. Organisms reproduce in ways specific to their life histories; communities may reassemble through dispersal and recruitment; soils and microbial communities mediate transformations in nutrients; hydrological processes connect sites across watersheds; and coastal or marine habitats depend on exchanges extending beyond administrative and industrial boundaries. These processes differ in timescale and cannot be represented by a single general rate of "regeneration."

Holling's foundational analysis of resilience distinguishes the persistence of system relationships under disturbance from narrow notions of local equilibrium stability [55]. Subsequent applications vary in what they call resilience, how they identify ecological states, and which disturbances are relevant. A

DRAFT

facility's recovery of electricity supply after an outage is consequently not evidence that the ecosystem supporting or receiving its material flows has recovered from its disturbances. Similarly, a count of organisms may return toward a recorded level while species composition, age structure, genetic diversity, food-web relationships, or habitat condition remain altered. Recovery of one measurable attribute must not be assumed to represent restoration of an entire ecological assemblage.

In a production-oriented inquiry, ecological continuation has at least four dimensions that should be examined separately. First, physical and biogeochemical processes remain within ranges relevant to the organisms and places being studied. Second, populations, communities, or habitats retain capacities for persistence and recruitment under the disturbances included in the analysis. Third, alternative trajectories, including adaptation and recolonization, are not foreclosed by poorly documented interventions. Fourth, the measurement and authorization of intervention respect appropriately articulated nonhuman interests and human communities affected by ecological change. The fourth consideration is normative and institutional; it should not be smuggled into the first three as if it were an ecological measurement.

The distinction also prevents a simplistic notion of technological substitution. The replacement of a damaged component may restore machine function, but purchasing machinery or planting a number of trees cannot automatically reproduce an old-growth forest's structure, a wetland's hydrological relations, or a long-lived marine habitat's ecological history. The possibility of restoration is a site-specific empirical and ethical matter. Section 22 will address recoverability and the limits of functional substitution in greater detail.

## 21.3 Contemporary Evidence of Ecological Pressures

Recent assessments underline the scale and interdependence of contemporary pressures. The *Planetary Health Check 2026*, published by the Potsdam Institute for Climate Impact Research in September 2026, reports that seven of nine assessed planetary boundaries are transgressed, including climate change, biosphere integrity, land-system change, freshwater change, biogeochemical flows, novel entities, and ocean acidification [154]. This is an integrated Earth-system assessment using control variables assembled from underlying observations and models, some predating its publication year. Exceedance of a planetary-boundary control variable is neither a direct observation of an ecological tipping event nor evidence that every local ecosystem has crossed the same limit. Translating global boundary assessments into site-specific decisions requires distinct ecological modelling, monitoring, and public or institutional judgment. In a distinct peer-reviewed assessment, Rockström and colleagues proposed safe and just Earth-system boundaries using global and subglobal criteria, explicitly distinguishing risks to Earth-system stability from significant human harm [159]. These proposals are informative research benchmarks; a production facility's legal or ecological acceptability cannot be read directly from a global boundary value.

The IPBES Global Assessment reported in 2019 that roughly one million animal and plant species were threatened with extinction, and assessed extensive human alteration of terrestrial and marine environments [59]. These figures characterize broad patterns of biodiversity decline, not experimentally isolated effects of AI or industrial robots. The IPBES Nexus Assessment, approved in December 2024, further assessed interdependencies among biodiversity, water, food, health, and climate and emphasized how isolated decisions can produce trade-offs across these domains [61]. The Global Resources Outlook 2024 documents the historical expansion of material resource extraction and the associated environmental burdens of economic activity [184]. Such assessments provide a reason to investigate material production carefully, while leaving open whether any particular automation pathway reduces, redistributes, or intensifies these pressures.

Table 33 distinguishes evidence of planetary pressures from the inference required for a particular industrial site. A high-level assessment can motivate hypotheses about freshwater withdrawals or land conversion without identifying the marginal impact of a new autonomous factory. Conversely, a small facility-level impact estimate does not settle its contribution to cumulative regional burdens. Both directions of inference need independently specified baselines, spatial boundaries, timescales, and causal

DRAFT

Table 33: Selected ecological assessments and the limits of their application to autonomous production. These sources do not estimate AI-production-specific ecological effects.

| Assessment | Reported finding | Scale and reference | Necessary qualification |
|---|---|---|---|
| Planetary Health Check 2026 [154] | Seven of nine planetary boundaries assessed as transgressed | Earth system; 2026 assessment of data from multiple periods | No facility-specific attribution, permit decision, or common local tipping point |
| IPBES Global Assessment [59] | Roughly one million animal and plant species threatened with extinction | Global synthesis published 2019 | Estimated threat, not a count of losses attributable to automation |
| IPBES Nexus Assessment [61] | Interdependencies and trade-offs among biodiversity, water, food, health, and climate | Global synthesis approved 2024 | Does not establish the outcome of a specified industrial intervention |
| Safe and just Earth-system boundaries [159] | Distinct proposed biophysical and significant-harm criteria | Global and subglobal modelling and literature synthesis, 2023 | Proposed boundaries are not universal local legal thresholds |

pathways.

## 21.4 Representing Coupled Production and Ecological Dynamics

The relational production structure of Section 6 can be extended through a typed description of ecological state and exposure. Let $p_n$ denote a vector of observable technical production states at discrete evaluation interval $n$; let $e_n$ denote a vector of ecological state variables defined for one or more specified places; and let $u_n$ denote a vector of production actions such as output scheduling, water abstraction, thermal discharge, land occupation, or waste-treatment decisions. Disturbances $\omega_n$ may include seasonal hydrology, extreme events, and ecological changes not caused by the facility. A provisional model is

$$\begin{aligned} p_{n+1} &= F_n(p_n, u_n, s_n, \omega_n), \\ e_{n+1} &= G_n(e_n, \mathcal{E}_n(u_{0:n}, s_{0:n}), \omega_n), \end{aligned} \tag{53}$$

where $s_n$ represents recorded upstream supplies and $\mathcal{E}_n$ is an exposure operator linking facility and supply-chain actions to ecological receptors. Equation (53) is a hypothesis class, not a calibrated ecological law. Its components may include different spatial and temporal resolutions; not every ecological effect can be reconstructed from the facility's own records. If delays or histories matter, the transition representation must include suitable state augmentation or memory, and if spatial heterogeneity is important it may require spatially explicit models rather than a small vector.

The components of $e_n$ require explicit scientific definitions and units. Illustrative candidates include water-table depth (metres), dissolved-oxygen concentration (mass per volume), habitat extent (area), population abundances (individuals under a stated sampling method), or concentrations of specified contaminants (mass per volume). Water withdrawals are measured as volumes per time; thermal discharge can be expressed as heat flux or an appropriate temperature exposure; land transformation is measured over area and duration. None of these observables should be summed in native units. Aggregation requires either a justified model connecting them to a common ecological endpoint or separate constraints evaluated as a vector. Monetary valuation, where relevant, constitutes an additional accounting layer and cannot replace the underlying ecological observables.

An illustrative ecological admissibility condition is

$$\mathcal{U}^{\text{ecol}}_{B,H} = \left\{ u_{0:H-1} : \Pr_{\mathcal{M}}[\, g_j(e_n, \ell_j) \le 0 \text{ for all } n \in \{0, \dots, H\},\ j \in J \,|\, \mathcal{I}_0] \ge 1 - \alpha \right\}, \tag{54}$$

where $B$ is the declared assessment boundary, $H$ the evaluation horizon, $\ell_j$ the location or receptor associated with constraint $j$, $\mathcal{M}$ a specified and assessable uncertainty model, $\mathcal{I}_0$ available initial information,

DRAFT

and $\alpha$ a declared tolerated probability of modeled constraint violation. The functions $g_j$ may differ in type and should be set using ecological knowledge and legitimate institutional procedures. A chance constraint is not a measure of moral permissibility. Even if Equation (54) holds under a simulation model, the inference may fail when omitted ecological processes, correlated disturbances, or incorrect observation models are significant. Where uncertainty cannot credibly be probabilized, robust scenario constraints or explicit uncertainty envelopes may be more appropriate.

A simple counterexample clarifies the separation at issue. For a constructed system, let technical output remain constant over a horizon while an ecological stock follows $e_{n+1} = e_n - \delta$, with $\delta > 0$ measured in consistent stock units per interval. The factory can satisfy a finite-horizon output target using existing equipment and supplies even though $e_n = e_0 - n\delta$ eventually crosses a specified ecological lower bound. The mathematical construction is elementary and is *not* evidence that every real factory depletes an ecological stock monotonically. Its role is to show why verifying technical output or uptime alone cannot establish ecological continuation. Conversely, ecological continuation at one measured receptor does not prove technical feasibility or wider ecosystem health.

## 21.5 Spatial Distribution and the Historical Relocation of Ecological Burdens

The dependency-relocation analysis in Section 9 extends naturally to ecological receptors. A factory may reduce on-site water consumption by obtaining upstream processed inputs whose manufacture is water-intensive. A more energy-efficient computational controller may rely on additional semiconductors whose fabrication, mineral refining, transport, and end-of-life treatment create burdens elsewhere. Material-recovery decisions may reduce primary extraction while increasing chemical-processing or transport demands in a different region. These are possible mechanisms, not general findings that autonomy necessarily causes ecological displacement. The empirical question is whether the new pathway's total and place-specific consequences are better, worse, or mixed under a matched production function and time horizon.

Industrial ecology, environmentally extended input–output models, material-flow analysis, and life-cycle assessment already provide relevant tools for tracing such pathways [184, 84, 83, 173]. Their interpretation requires care. National or sectoral averages may not characterize the ecological exposure at a particular river basin, refinery, mine, or habitat. A global-warming-potential indicator captures a defined climate forcing comparison but does not identify effects of local chemical toxicity, freshwater depletion, acoustic disturbance, or changes in community composition. Allocating emissions or materials among co-products requires declared rules and sensitivity analyses. A GR trajectory record is useful only to the extent that its historical, institutional, and provenance fields improve traceability or causal discrimination beyond a competent conventional inventory.

Spatial distribution also creates asymmetries of observability and authority. The enterprise selecting an automation strategy may observe product yield and on-site energy consumption precisely while possessing only modelled estimates of contamination or land change at distant upstream sites. Regulators and nearby communities may have different measurements and different rights to challenge the adequacy of the assessment. Thus, an ecological dependency is not exhausted by the presence of a material-flow edge: the availability, reliability, circulation, and contestability of evidence affect which ecological consequences become visible in production decisions. The paper's GR contribution is to retain these heterogeneous relations explicitly without conflating an epistemic power to select measurements with a biophysical mechanism of ecological change.

## 21.6 Ecological Timescales and the Limits of Technical Substitution

Technical assets, material stocks, and ecosystems can exhibit sharply different characteristic timescales. Controllers may be replaced within weeks, facilities refurbished over years, and some ecological relationships altered or reassembled over far longer periods. Recolonization does not necessarily restore a prior food web, and the renewal of water stocks does not by itself guarantee the recovery of aquatic communities. Relevant disturbances may also be cumulative: intermittent extraction, repeated thermal

DRAFT

releases, landscape fragmentation, and combined contaminants can interact even when each operation passes a short-term compliance check. Holling's resilience framework and the IPBES Nexus Assessment motivate treating such interactions as specific ecological research problems rather than assuming linear recovery after cessation of a disturbance [55, 61].

The choice of evaluation horizon therefore has normative and epistemic consequences. A five-year engineering horizon may accurately measure replacement costs yet exclude effects on long-lived organisms or slowly changing water and soil systems. Extending the horizon is not a universal remedy: observations become sparse, counterfactual assumptions less testable, and ecological change may involve multiple possible pathways. A responsible analysis specifies which future trajectories are being assessed and the uncertainty surrounding their continuation, rather than treating ecological states at designated dates as guaranteed outcomes. It should also consider whether some losses cannot be meaningfully compensated through restoration elsewhere or through monetary payments.

This is a point at which ecological ethics and historical materialism meet without being collapsed. A particular mode of production determines how investment decisions, material access, risk allocation, and ecological monitoring are organized. But the ecological consequences are mediated by physical processes not constituted by ownership alone. Changing a firm's ownership may change the incentives and available remedies while leaving a contaminant's biochemistry unchanged. Conversely, technical innovations can alter exposure mechanisms, yet their selection and scale depend on social and institutional relations. No ownership form or technology is assumed to produce ecological compatibility intrinsically.

## 21.7 Nonhuman Interests and the Justification of Ecological Intervention

Ecological assessments often begin with processes valuable to humans: the renewal of water supplies, pollination, climate regulation, food production, or protection against hazards. These matters are important, but an adequate account should also ask how productive interventions affect organisms and ecological assemblages whose continuation may matter independently of industrial utility. An ecosystem-level indicator can miss the exposure of particular animals; a species-abundance indicator can obscure individual welfare; and a short-term biological response cannot establish the long-term reproductive consequences of repeated interventions. Nor should the presence of ecological uncertainty automatically be converted into an unconditional prohibition on productive activity. The ethical and institutional task is to articulate relevant interests, uncertainty, decision authority, alternatives, and possibilities for avoiding or remedying harm.

The distinction between *measurement*, *interpretation*, and *authorization* is decisive. A sensor may establish a dissolved-oxygen concentration or heat flux. An ecological model may suggest an altered habitat risk. A legal regime or governance institution may determine which interventions can proceed and which monitoring or remedial requirements apply. These inferential and institutional steps can be mutually informing, but none substitutes for another. Planetary-boundary assessments, for example, offer a scientifically structured view of systemic risks; they are not themselves permits, universal local thresholds, or adjudications of nonhuman rights [159, 154]. This manuscript does not posit that a single formal evaluator can express all human and nonhuman interests in a common scalar welfare function.

A historical-materialist comparison should consequently examine how capitalist enterprises, public facilities, cooperatives, and commons-oriented infrastructures define ecological obligations, obtain evidence, share decisions, and allocate the capacity to contest interventions. Ecological compatibility could improve in more than one arrangement, and different dimensions might diverge. The analysis should admit findings in which a highly automated private facility performs better on declared ecological indicators than a publicly operated comparison, or vice versa. The proposed inquiry concerns the conditions and causal mechanisms underlying such differences, not the presumed virtue of an institutional label.

## 21.8 Empirical Tests of Ecological Productive Continuation

A testable comparison could examine alternative automation, sourcing, and cooling policies for a defined productive service, such as a specified volume and quality of manufactured components over a fixed

DRAFT

Table 34: Evidence interfaces for evaluating ecological continuation under specified production policies. The proposed measures are complementary, not interchangeable indicators of sustainability.

| Interface | Required observations | Principal uncertainty | Invalid inference to avoid |
|---|---|---|---|
| Technical service | Qualified output, uptime, repair and input stocks | Demand and equipment failures | Stable output proves ecosystem viability |
| Physical exchange | Water, energy, material, heat and chemical flows | Allocation, off-site sources and losses | Lower on-site flow proves lower total burden |
| Ecological receptors | Habitat and community indicators, local exposures, reference variation | Detection bias, delayed and interacting responses | One recovered metric proves ecosystem restoration |
| Spatial supply chain | Location-specific sourcing and processing, impacts at critical nodes | Missing tiers and sectoral-average substitutions | National intensity captures every local exposure |
| Institutional evaluation | Monitoring provenance, participation, authorization and remedies | Incomplete rights or undocumented decisions | A modelled boundary confers legal permission |

operating horizon. The direct measurement programme would record equipment performance, electricity, material inputs, water withdrawal and consumption, heat rejection, process chemicals, losses, and on-site environmental conditions. The indirect programme would estimate or measure major supplier-stage burdens while identifying which upstream links rely on averages rather than facility data. Independent ecological observation would specify receptors, baseline variability, comparison areas, habitat processes, sampling frequency, and uncertainty. Community and institutional records would identify the procedures through which monitoring claims and intervention decisions are evaluated and challenged.

Table 34 sets out interfaces to be connected without collapsing their evidential standards. Three comparisons would be particularly informative. The first would hold product specifications fixed while comparing conventional and AI-assisted control policies using equivalent ecological endpoints. The second would compare on-site improvements with upstream exposures under alternative sourcing arrangements. The third would compare different institutional monitoring and access arrangements while retaining the same physical technology, where a credible comparison is available. Each comparison would include prospective stopping criteria for unacceptable or poorly characterized ecological impacts rather than requiring disturbance to continue for the sake of a complete experiment.

Several negative results would be theoretically useful: no measurable difference in ecological exposure under matched policies; improvements in engineering efficiency offset by increased output volumes; superior on-site performance accompanied by greater upstream burdens; or a richer GR relational record that fails to improve environmental attribution relative to standard life-cycle and provenance tools. Stronger positive findings would require appropriately calibrated models, adequate controls, reproducible measurement, and evidence of ecological outcomes across the relevant time horizons. The mere number of recorded relations or the computational complexity of a controller is not an ecological endpoint.

The section thus establishes a constrained conclusion. The reproduction of productive infrastructure takes place through material relations embedded in ecological systems that evolve according to heterogeneous processes and histories. Technical continuity, even when supported by renewable energy, sophisticated control, maintenance, and recovery, does not itself determine whether these processes can continue within declared ecological and normative limits. The next section examines ecological recoverability more closely, distinguishing measurable improvement, re-establishment of particular ecological functions, and the limits of restoration following productive intervention. The subsequent ecological sections will test how such conditions can be integrated into lifecycle accounting and robust viability

DRAFT

analysis without presuming that an autonomous production system can fully know, control, or replace its ecological surroundings.

# 22 Ecological Recoverability Following Productive Intervention

Section 21 distinguished the continued operation of a productive installation from the continuation of ecological processes exposed to its material activities. This section asks a further question: after a specified productive intervention changes an ecological system, which aspects of that system can recover, through which processes, and with what evidential and institutional qualifications? The research purpose is to establish a defensible account of *ecological recoverability* before subsequent sections consider lifecycle impacts, robust ecological viability, and cumulative changes at scale. The analysis draws upon restoration ecology, resilience research, empirical studies of disturbed habitats, and historical materialism. It treats recovery as a set of observable and potentially contested processes rather than a presumed return to an unchanging natural baseline. The account is prospective and comparative; it supplies neither a universally applicable recovery threshold nor an empirical ecological assessment of a particular autonomous factory.

## 22.1 Specifying Recovery in a Changing Ecological System

Ecological recovery is ordinarily evaluated relative to attributes of ecosystems, such as species composition, structural complexity, hydrological functioning, nutrient cycling, reproductive recruitment, and the capacity to respond to subsequent disturbance. These attributes can develop along different trajectories. A polluted freshwater reach, for example, may show rapidly declining contaminant concentrations even while sensitive invertebrates recover slowly; vegetation cover can return while characteristic soil communities remain altered; and regenerated biomass can coexist with a reduction in ecological connectivity. An explicit statement of what is expected to recover is therefore logically prior to a declaration of success.

The Society for Ecological Restoration's third-edition international standards distinguish ecological restoration directed toward native ecosystems from the broader restorative continuum, emphasize reference models that account for environmental change, and call for measurable evaluation of recovery attributes [46]. These principles help avoid two symmetrical errors. A reference state should not be treated as a timeless or universally attainable configuration, while an arbitrary improvement in one indicator should not suffice to establish the recovery of the ecological relations relevant to the stated goal. Reference models can draw on multiple relatively intact sites, historical observations, Indigenous and local knowledge where appropriately governed, and defensible reconstructions of ecological potential. Their selection carries scientific and institutional responsibilities, including uncertainty regarding changing climate, disturbance regimes, and historical land use.

For this paper, let $\mathcal{R}_{x,H}$ denote an explicitly described set of acceptable ecological reference trajectories for location or ecological unit $x$ over a declared horizon $H$. The set may contain multiple trajectories where climate variability, succession, or competing plausible reconstructions prevent one uniquely specified future. Such a representation differs from a single fixed target: the question is whether a proposed recovery process remains compatible with a declared range of ecological attributes and continuations. The construction of $\mathcal{R}_{x,H}$ requires substantive ecological judgment and cannot be inferred from the notation itself. A reference trajectory can also be normatively inadequate if it ignores threatened organisms, nearby communities, or obligations to affected people.

A useful vocabulary separates *attenuation of exposure*, *biological re-establishment*, *recovery of measured functions*, and *re-establishment of continued ecological capacities*. Exposure can decline following emission cessation. Recruitment and recolonization may occur if propagules, habitats, and connectivity are available. Selected processes may recover without a corresponding return of previous community composition. Finally, durable ecological continuation depends partly on the system's capacity to organize and renew its own relations under subsequent environmental variability. These distinctions specify

DRAFT

empirical research tasks; they do not assume that recovery follows a universal sequence.

## 22.2 What Restoration Evidence Establishes

Empirical restoration studies show both substantial possibilities for biological recovery and important limits to its timing and completeness. A global meta-analysis by Moreno-Mateos and colleagues assembled evidence from 621 restored or created wetland sites. On average, measures of biological structure and biogeochemical functioning remained approximately 26% and 23% lower than reference wetlands, respectively, even in a synthesis including sites assessed after approximately a century of restoration [129]. Those percentages are comparisons within the study's wetland sample and selected metrics. They are not universal deficits for every restored wetland, nor direct predictions of effects at an industrial production site. The study further identified differences associated with wetland size, climate, hydrological exchange, and environmental setting. The important inference concerns heterogeneity: the measurement of an improving attribute does not settle whether the wider set of ecological attributes has recovered.

A meta-analysis of 133 tropical-forest studies by Crouzeilles and colleagues compared natural regeneration with active restoration, controlling for selected environmental and historical factors. Across the study's biodiversity and vegetation-structure measures, natural regeneration performed better on average than active restoration [25]. This finding cautions against equating the degree of technical intervention with ecological restoration success. It does not imply that natural regeneration without further site-specific management is superior in every degraded forest, or that withholding necessary intervention is always appropriate. Different disturbance legacies, proximity to seed sources, invasive species, land-use pressures, and reference conditions may justify different strategies. An autonomous production system's restoration plan must similarly explain why a proposed intervention helps remove limiting conditions rather than merely increasing the visible activity of restoration.

Marine evidence illustrates the possibility of contrasting recovery responses at the same site. Jones and colleagues evaluated the area of an experimental polymetallic-nodule mining track 44 years after a 1979 test. They observed continuing physical alteration and persistent biological differences in directly disturbed areas, while several organism groups had begun to re-establish; some areas affected by sediment plumes showed more limited detectable residual effects [108]. The study concerns an intensive historical deep-sea mining test with particular equipment, substrate, and habitat conditions. It cannot directly estimate the disturbance of ordinary marine research robots, nor the recovery rate of terrestrial autonomous manufacturing infrastructure. Its methodological relevance lies in the separation of physical persistence, biological recolonization, and response differences among exposure zones and organism groups.

Table 35 records a heterogeneous evidence base instead of converting its findings into an average recovery percentage. Such conversion would confound habitats, disturbance intensities, observations, and reference constructions. A historically situated analysis also asks which activities created the degradation, who can access the affected environment, which restoration knowledge is recognized, and whether restoration funding and monitoring can continue after the production activity ends.

## 22.3 Ecological Memory and Recovery Constraints

Ecological recovery is conditioned by more than the contemporaneous intensity of a disturbance. Prior land use may change soil profiles and seed banks. Industrial extraction may remove physical substrates used by sessile organisms. Contaminant residues may persist in sediments, while groundwater transport can delay responses after surface emissions cease. Fragmented habitats may lack connected populations able to recolonize a site. Similarly, repeated heat or drought stress can interrupt biological recruitment even where a short-term indicator appears favourable. These mechanisms introduce *ecological memory*: the consequences of earlier exposures and ecological relations that remain causally relevant to later outcomes.

Resilience theory has long distinguished the persistence of selected ecosystem functions from the capacity of systems to absorb disturbance while retaining relevant organizing relations [55]. This liter-

DRAFT

Table 35: Selected empirical evidence on ecological recovery and the limits of transfer to autonomous production. Dates, populations, and metrics refer to the cited studies rather than a shared global recovery index.

| Evidence setting | Observation design | Relevant finding | Boundary of inference |
|---|---|---|---|
| Restored wetlands [129] | Meta-analysis of 621 sites across multiple regions | Average biological structure and biogeochemical functions remained below reference measures, including in long-duration comparisons | Habitat-, metric-, and study-design dependent; no universal recovery rate |
| Tropical forests [25] | Meta-analysis of 133 studies | Natural regeneration exceeded active restoration on studied biodiversity and vegetation-structure outcomes, conditional on included controls | Average across sampled forest settings; no presumption against necessary active intervention |
| Abyssal mining track [108] | Field re-examination 44 years after a 1979 mining test | Persistent physical alteration and taxon-dependent biological re-establishment | Intensive mining disturbance, not routine industrial or research robotic operation |
| Restoration standards [46] | Expert-developed normative and practice guidance, third edition (2026) | Reference models, recovery processes, monitoring, and affected-party involvement are central to restoration practice | Practice guidance rather than a measured causal effect |

ature provides an intellectual precedent for the paper's provisional GR interest in the continuation and transformation of relational capacities. The formalism should add value by making heterogeneous dependencies, exposed organisms, and contested reference choices traceable, without claiming to displace established ecosystem dynamics, resilience measures, or ecological network models. A recorded history is useful only to the extent that it improves causal attribution, explanation, comparison, or the practical capacity to identify and correct harmful interventions.

An important constraint is the possible presence of nonlinear responses and path dependence. An ecosystem driven beyond a threshold may occupy a different configuration even after the original pressure decreases. The existence and location of such thresholds are habitat-specific empirical questions. They should not be inferred merely from the mathematical convenience of representing alternative stable states, nor presumed absent because no abrupt transition has yet been observed. Recovery evidence may be weak at the temporal scale of a production investment cycle, particularly for slow-growing organisms, deep sediments, groundwater systems, or processes whose observable responses are delayed.

The relationship to productive reproduction is direct. Maintaining industrial equipment can often draw upon replaceable components, standardized interfaces, and qualified repair procedures. Ecological processes may involve historically contingent organisms, species interactions, substrates, and landscapes for which no materially equivalent manufactured replacement exists. A production plan that presumes future ecological repair therefore carries a separate burden of evidence for the availability of its ecological recovery pathways, rather than borrowing confidence from the repairability of the machines conducting the intervention.

## 22.4 Representing Recovery as a Bounded Possibility

A minimal formal model can represent recovery without asserting that ecological systems are fully observed or controllable. Let $e_n \in \mathcal{E}$ denote an ecologically meaningful, generally partially observed state at observational index $n$. Let $u_n \in \mathcal{U}(e_n)$ represent admissible restorative or non-intervention decisions, $\omega_n$ an uncertain physical and biological disturbance, and $\theta$ a collection of model parameters. We write

$$e_{n+1} = G_n(e_n, u_n, \omega_n; \theta), \tag{55}$$

where admissibility of $u_n$ includes declared ecological and institutional restrictions. The time index designates observation and decision occasions, rather than presuming the ontology of ecological time. In practice, $G_n$ would need habitat-specific evidence and uncertainty models.

For a declared horizon of $H$ observation intervals, let $\mathcal{K}_H(x) \subseteq \mathcal{E}$ denote the ecologically acceptable continuation states derived from the reference model, endpoint definitions, and normative review for site $x$. Let $\Omega_H$ collect the admitted disturbance trajectories. We define a provisional *robustly recoverable set*

$$\mathcal{R}_H^{\text{rob}}(x) = \left\{ e_0 \in \mathcal{E} : \ \exists \pi \in \Pi_{\text{adm}} \, \forall \omega_{0:H-1} \in \Omega_H, \ e_H^{\pi,\omega} \in \mathcal{K}_H(x) \right\}. \tag{56}$$

Here $\pi$ is a nonanticipative state- or observation-dependent policy satisfying the specified admissibility requirements. Its available observations and whether the initial state is known must be declared when defining $\Pi_{\text{adm}}$. Otherwise a policy selected separately with knowledge of each hidden $e_0$ could create a spurious recovery guarantee under partial observation. Equation (56) addresses a terminal continuation condition; more stringent evaluations may also require acceptable conditions at intermediate states or persistence after $H$. Replacing the universal disturbance requirement with probabilistic confidence would define a different claim. If disturbances are unbounded, targets too restrictive, or restoration choices insufficient, the robustly recoverable set may be empty. Thus the definition specifies the evidence a claim would require, without proving that restoration is feasible.

There is also a more basic identification limit. Suppose the actual ecological state is $e = (a, b)$, where $a$ denotes an observed attribute such as visible vegetation cover and $b$ an unobserved attribute such as a relevant soil-community composition. If the observation mechanism is $h(e) = a$, then distinct states with the same $a$ are observationally indistinguishable. This leads to the following elementary proposition.

**Proposition 13** (Non-identification of full recovery from a non-injective observation)**.** *Let $h : \mathcal{E} \to \mathcal{Y}$ be a measurement map, and let $\mathcal{K} \subsetneq \mathcal{E}$ be the specified set of ecologically recovered states. If there exist $e^{\text{in}} \in \mathcal{K}$ and $e^{\text{out}} \notin \mathcal{K}$ such that $h(e^{\text{in}}) = h(e^{\text{out}})$, then observing $h(e)$ alone does not identify whether $e \in \mathcal{K}$.*

*Proof.* The two states induce the same observation, but one satisfies the recovery criterion and the other does not. An inference using only that observation gives the same information for both and therefore cannot distinguish membership in $\mathcal{K}$. □

The proposition is an elementary identification statement. The failure of injectivity on the entire state space is not itself sufficient to prevent identification of recovery membership: what matters is whether any observation-equivalence class intersects both $\mathcal{K}$ and its complement. Conversely, full-state identification may be unnecessary when observations and warranted restrictions separate those two classes. The proposition establishes no empirical irreversibility or novel ecological law. Its relevance is that a recovery certificate based on a limited observable requires justification that relevant unobserved states cannot invalidate the claimed recovery, or must carry an explicit statement of its restricted endpoint scope. Sampling additional independent attributes, using ecological experiments, and constraining the plausible state space can improve identification, but only under defensible measurement and ecological assumptions.

## 22.5 Limits of Substitution and Ecological Compensation

The recovery of ecological capacities poses distinct questions from economic replacement of manufactured inputs. A production operator may be able to substitute a different component to restore technical output. An ecological substitute would have to reproduce the relevant processes and relations at the affected site or across an explicitly justified ecological unit, including where appropriate habitat structure, species assemblages, connectivity, genetic diversity, nutrient cycling, and the continuations of nonhuman lives. Some characteristics can be enhanced through engineering, and restoration can achieve substantial benefits. Other losses may be poorly understood, geographically non-transferable, or effectively irreversible within a declared horizon.

This matters for claims of environmental offsets and compensation. A reconstructed habitat elsewhere can have measurable conservation benefits, yet benefits at a different location or time need not

DRAFT

Table 36: A prospective research design for evaluating recovery claims. Each row identifies an outcome requiring its own observations and inferential qualifications.

| Research interface | Required measurements or records | Comparison strategy | Claim not established by the indicator alone |
|---|---|---|---|
| Disturbance and exposure | Baseline conditions, area and intensity of alteration, contaminant/heat/water pathways | Matched sites and pre/post measurements where feasible | Reduced exposure equals recovered ecosystem |
| Biotic re-establishment | Recruitment, species composition, traits, genetic or population indicators | Repeated surveys with detection correction | Higher abundance restores prior relations |
| Ecological function | Hydrology, nutrient cycling, soil/sediment processes, productivity | Appropriate reference ranges and process assays | One functional metric captures ecosystem integrity |
| Persistence and resilience | Repeat disturbance, connectivity, longer monitoring horizon | Site histories and stress-response comparisons | Short-term endpoint implies durable continuation |
| Governance and remedy | Reference-choice record, affected-party involvement, obligations, funding, correction mechanisms | Compare decisions and accountability over time | Documentation proves legitimate authorization |

restore affected organisms, cultural relations, or ecological dependencies at the damaged site. An accounting scheme that credits restored area or predicted service delivery against measured habitat loss therefore requires explicit ecological equivalence, uncertainty, timing, and distributional assumptions. Legal acceptability and ethical justification constitute further questions. It would be misleading to treat a positive monetary valuation, the planting of replacement vegetation, or a contracted restoration budget as an observed recovery outcome.

A historically informed analysis should also examine the allocation of recovery obligations. Productive expansion can generate immediate revenue while its ecological effects persist beyond a facility's ownership period or the lifespan of the equipment used. Monitoring, liability, environmental remediation, and restoration funding therefore have temporal and institutional dimensions. The consequences of an intervention may remain with affected ecosystems and communities after economic beneficiaries, controlling entities, or technologies have changed. Autonomous decision systems make provenance and accountability particularly relevant, because recommendations, human approvals, procurement choices, and control actions may be distributed across organizations. The GR representation may assist in tracing these relations; it cannot, by recording them, confer legitimacy or substitute for environmental remedy.

## 22.6 Comparing Recovery Pathways under Productive Constraints

A credible empirical comparison requires a concrete site, specified disturbance, realistic alternatives, and repeated observations. Consider a hypothetical industrial expansion adjacent to a wetland. Three possible strategies might be avoidance of the most sensitive area, minimized construction with compensatory on-site rehabilitation, and greater construction disturbance followed by more extensive active restoration. Their engineering outputs and ecological consequences may differ; comparison must declare production-service equivalence, baseline variation, affected habitat, social impacts, funding reliability, and uncertainty. The purpose is to test rather than presuppose the comparative success of the more technologically intensive option.

Where practical, a before–after–control–impact design could improve attribution, although spatially replicated controls and treatment assignment may be difficult and causal interpretation remains dependent on parallel-trend or site-comparability assumptions. A control area should not simply be selected because it is currently less disturbed; its history, physical setting, and ecological connectivity matter. Studies

DRAFT

should also record climate and background stressors, evaluate monitoring sensitivity, and specify the interval over which a change could reasonably appear. Short-term recovery of a visible endpoint should not be treated as evidence of long-term ecosystem resilience.

Table 36 provides a matrix for testing claims made by autonomous production systems. An AI controller could select lower-impact construction routes, optimize energy or water demands, or adaptively adjust remediation. Evidence of improved planning quality would be a relevant engineering finding. The ecological claim would require measured changes in exposures and outcomes, with independent checks for unanticipated disturbance and omitted receptors. Negative findings remain informative: an optimized policy may yield no detectable ecological gain, a naturally regenerating site may outperform an actively restored site, or costs and ecological burdens of restoration may exceed the stated benefits under particular boundaries.

## 22.7 Recoverability within Historical Production Relations

Historical materialism helps explain how different production relations shape exposure, the resources available for rehabilitation, and the distribution of responsibility. A production system's ecological decisions arise within procurement contracts, property and land-use institutions, insurance and liability arrangements, labour structures, and public or collective mechanisms for evaluating environmental interventions. These relations affect the choice of sites, the time horizons of investment, the visibility of ecological harms, and the capacity to fund monitoring after extraction or facility closure. They do not determine biological recovery in isolation; ecological responses depend on the physical disturbance and the historical properties of the affected system.

The distinction between productive reproduction and ecological recoverability sharpens the paper's political-economic hypothesis. Automation may increase the ability of enterprises to renew productive equipment while allowing ecologically costly effects to be shifted outside financial reporting boundaries or into longer time horizons. Under other institutional and technical circumstances, improved control, lower material throughput, monitoring, and enforceable restoration obligations may reduce impacts or support ecological recovery. A comparative study must allow both possibilities. These possibilities should be investigated through matched service outputs, physical material flows, independently monitored ecological attributes, and documented institutional decision rights.

Ecological continuation also cannot be reduced to the productivity of nature for an enterprise. Nonhuman organisms and ecological relations may have claims to ethical consideration irrespective of immediate use-value. Restoration goals involve choices about which beings and trajectories deserve protection, which losses are tolerable, and who can authorize action under uncertainty. The research model can make those choices more legible, but it supplies no automatic solution to conflicts among different ecological and human interests. In particular, a facility's capacity to pay for environmental rehabilitation neither demonstrates that the original intervention was justified nor assures that the required ecological recovery will occur.

## 22.8 Implications for Ecological Viability Research

The resulting conclusion is intentionally bounded. Ecological recoverability is contingent on the nature of the disturbance, the ecological attributes considered, available recolonization and renewal processes, the sequence of interventions, background environmental changes, and the evidential and institutional standards used to evaluate outcomes. A system that restores vegetation cover, lowers an emission rate, or increases the observed abundance of a selected taxon may have achieved an important improvement, while wider claims about the return of ecosystem integrity or generative capacity require additional evidence. Conversely, some systems recover through processes that do not reproduce their former measurable configurations exactly, and the relevant normative question may concern the continuation of defensibly identified ecological relations rather than historical replication.

For AI-mediated productive systems, the research implication is that prospective recovery should not be credited as an assured capacity simply because a restoration policy can be represented, simu-

DRAFT

lated, or financed. Before invoking recovery to justify productive expansion, a responsible assessment would specify the ecological state variables, reference alternatives, experimental or observational basis, uncertainty, time horizon, material limits, and accountable institutions. The formal recoverable set in Equation (56) is a research object whose membership requires habitat-specific investigation, rather than an instrument that can certify a proposed project by itself.

Section 23 will place the direct disturbance, recovery activities, upstream supply chains, and equipment lifecycles within compatible accounting boundaries. Section 24 will then investigate which production policies can satisfy ecological constraints under declared uncertainty. Section 25 will consider the cumulative consequences when many facilities and repeated interventions interact across space and time. This order preserves the distinction among observed ecological responses, lifecycle burdens, formal viability, and institutional judgments, and makes room for negative evidence at every interface.

# 23 Accounting for Ecological Burdens across Production Lifecycles

Sections 21 and 22 established why the operation of industrial equipment and the recoverability of an exposed ecological system require different observations. The present section examines a further methodological problem: how can the ecological burdens associated with an AI-mediated productive service be traced through the equipment, materials, energy, maintenance, distribution, and recovery activities on which that service depends? Its role is to connect previously developed dependency models to established life-cycle accounting, while preserving the distinction between inventories of physical exchanges, models of potential ecological impact, observed ecological change, and judgments concerning acceptable intervention. The discussion draws on life-cycle assessment (LCA), environmentally extended input–output analysis, material-flow methods, and emerging spatially differentiated impact assessment. The proposed generative-relational (GR) representation is intended to make changing dependencies and decision rights more traceable; it is not offered as a replacement for competent industrial-ecological accounting or as an independently validated impact model.

## 23.1 Specifying the Productive Service and Assessment Boundary

A comparison requires an explicit *functional unit*: the quantified service or function against which alternatives are evaluated. A robot-integrated production cell and a less automated cell, for example, cannot be compared meaningfully by their electricity demand per operating hour if they deliver different numbers or qualities of conforming products. A prospective functional unit might be the delivery of ten thousand conforming precision components to a stated specification, within a declared period and region, with recorded delivery reliability and inspection requirements. The example is a research-design choice, not a measured output of an existing installation. Comparability may require adjustment for rejects, rework, product durability, downstream use, and the number of facilities needed to supply the same service.

ISO 14040 specifies the framework of goal-and-scope definition, inventory analysis, impact assessment, interpretation, and reporting; ISO 14044 sets requirements and guidelines for carrying out the assessment [84, 83]. A declared *reference flow* then specifies what quantities of goods, services, and activities supply the functional unit. This distinction matters where the reference flow changes with equipment lifetime or production architecture. The production cell might require one controller and a replacement module over a five-year horizon, while its comparator uses different tooling and more frequent operator interventions. Assigning the entire manufacturing burden of a long-lived machine to one short production run would distort such a comparison; assigning it uniformly over an assumed service lifetime might also misrepresent irregular operation or premature failure. Utilization, service life, refurbishment, and replacement therefore require evidence and sensitivity analysis.

System boundaries should extend beyond the factory gate where relevant. At minimum, the manuscript proposes a *foreground* inventory of the production process, its direct energy and water inputs, consumables, maintenance, rejected products, and process emissions. A *background* inventory traces electricity generation, equipment fabrication, semiconductor supply, mining and refining, trans-

DRAFT

port, and waste treatment. Infrastructure construction and decommissioning must be included when consequential to the comparison. The boundaries also require a time horizon, geography, policy scenario, and definitions of which emissions or resource flows are attributable to the service. A nominally cradle-to-grave assessment does not demonstrate comprehensive coverage: missing upstream tiers, poorly understood land transformation, and untracked residual waste can persist within the declared boundary. The Product Environmental Footprint methods recommended by the European Commission provide one established framework for consistent life-cycle evaluation, while remaining methods whose relevance depends on product-specific rules and study purpose [33, 183]. Methodological standardization of an inventory does not itself supply unmeasured exposure pathways, receptor observations, or causal evidence of ecological recovery.

## 23.2 Accounting for Physical Exchanges without Collapsing Units

The inventory should preserve quantities by physical type. Kilowatt-hours of electricity, litres of withdrawn water, cubic metres of consumed water, kilograms of copper, hectares of habitat conversion, and kilograms of a named contaminant are not interchangeable measures of ecological burden. Nor should water withdrawal be silently equated with consumption, or the movement of collected scrap with the production of qualified secondary material. Sections 17 and 20 gave examples of these distinctions in energy, cooling, and recovery systems. The objective here is to connect those localized accounts to the complete service pathway without combining dimensions through an unexplained coefficient.

Let $a \in \mathcal{A}$ index productive or maintenance activities, $r \in \mathcal{R}$ regions or receptor locations, $n \in \mathcal{N}$ time intervals, and $k \in \mathcal{K}$ elementary flow types. A proposed inventory record is

$$I^{(\pi)}_{arnk} = \text{quantity of typed elementary flow } k \text{ attributable to activity } a \text{ under policy } \pi \text{ at } (r, n), \quad (57)$$

with an explicit unit $[I_k]$ for each $k$. Production events, purchases, and recovered materials should be identified separately from elementary exchanges with the environment so that intermediate inputs are not counted repeatedly. A provenance record would connect each entry to a measurement or estimator, its time stamp, location resolution, allocation rule, uncertainty, and the decisions that selected the underlying activity. In particular, the historical record of production instructions should not be mistaken for evidence that upstream exchanges have been measured.

For an impact category $j$, a conventional characterization-based approximation may take the form

$$\widehat{B}^{(\pi)}_j = \sum_{a,r,n,k} c_{jkrn}\, I^{(\pi)}_{arnk}, \qquad [c_{jkrn}] = [B_j]/[I_k], \quad (58)$$

when the model is linear in the relevant elementary flows and the characterization factors $c_{jkrn}$ are defensible for the context. These factors are not universal physical constants. They encode environmental fate, exposure, or impact assumptions under a selected model. Equation (58) describes a potential impact estimate in the units of category $j$, not a measured change in an entire ecological system. Some impacts exhibit nonlinear responses, thresholds, interactions, or time dependence not captured by a linear factor. For such cases the impact operator must be replaced or supplemented by a suitably calibrated ecological model, or the missing relationship reported as unknown.

The outputs should therefore be a *vector* of non-equivalent indicators: for example climate-change potential, freshwater deprivation risk, modeled toxicity potential, land-use pressure, and other pathway-specific ecological metrics. An index combining different categories requires declared normalization and weighting choices and should be reported alongside the original categories. The IPBES Values Assessment documents the diversity of ways in which nature and its contributions are understood and valued, reinforcing the distinction between an impact calculation and a comprehensive account of ecological or ethical value [60]. No single scalar LCA score will be treated here as an intrinsic measure of ecological generativity.

DRAFT

## 23.3 Tracing Upstream and Downstream Production Pathways

A production facility's direct inventory cannot describe the full material consequences of its productive policy. Section 21 established this point in the context of displaced electricity, materials, and ecological exposures. A robot-operated facility may use less electricity on site while depending on additional sensors, network infrastructure, accelerators, and replacement devices. Conversely, greater equipment durability or process control might reduce rejected materials, component turnover, or upstream extraction. Both changes must be counted in the same functional comparison. Increased operational autonomy alone determines neither outcome.

Process-based LCA can follow specific production steps with detailed engineering inventories, while environmentally extended multiregional input–output (EE-MRIO) approaches trace interindustry purchases and associated environmental extensions across large production networks. EXIOBASE 3 provides an influential framework of the latter type; its foundational publication describes a harmonized 1995–2011 time series with 44 explicitly represented countries, five residual world regions, and social and environmental satellite accounts [173]. These historical tables cannot be taken as a direct observation of a particular factory or of contemporary industrial supply conditions. The strength of MRIO accounting is broader upstream coverage; its weaknesses for a particular factory include aggregated sectors, average technology assumptions, region-level resolution, and uncertainty in mapping specialized products to economic sectors. Process inventories are often more specific but may truncate unobserved upstream tiers. Hybrid approaches can use the strengths of both, provided the overlap between process and input–output accounts is identified and deducted rather than double counted [111].

The standard Leontief accounting relation, under its familiar assumptions, provides an interpretive starting point:

$$\mathbf{x} = (\mathbf{I} - \mathbf{A})^{-1}\mathbf{y}, \tag{59}$$

where $\mathbf{A}$ represents monetary or appropriately constructed physical requirements, $\mathbf{y}$ final demand, and $\mathbf{x}$ total gross output. With compatible environmental-extension matrices, attributed upstream pressures can be estimated. This equation is established input–output mathematics, not a new theorem or a complete physical model of ecological interactions. If $\mathbf{A}$ is a monetary transactions matrix, its coefficients must not be interpreted as tonnes or megajoules without an explicit conversion. The inversion also assumes a fixed technical structure within the modeled comparison, which may be inadequate where automation, substitution, energy mixes, or trade routes change significantly. Sensitivity and scenarios are consequently part of the measurement architecture, not merely supplementary discussion.

Downstream stages require equal care. Maintenance and re-commissioning consume materials and energy throughout the assessment horizon, and equipment may have a second functional life under refurbishment or reuse. Recycling can create qualified substitute material while incurring collection, transport, chemical processing, and unrecovered residual burdens. Attributing a credit for avoided primary production requires a documented counterfactual substitution scenario rather than an unconditional subtraction. End-of-life and reuse pathways should be checked for quality differences, leakage, uncertainty, and effects on the actual supply of primary inputs. The GHG Protocol product standard addresses product lifecycle greenhouse-gas emissions and removals. Avoided emissions attributable to a comparison with an alternative product or service are a distinct counterfactual claim: they must not be deducted from the product inventory, although separately supported and disclosed estimates may be reported [195]. Accordingly, neither a smaller product GHG inventory nor a separately reported avoided-emissions estimate automatically establishes improvements in water, toxicity, biodiversity, or local habitat processes.

## 23.4 Allocation and the Counterfactual Meaning of Improvement

Autonomous production frequently generates more than one output: conforming products, recoverable heat, scrap, secondary materials, or reusable equipment. A unit of equipment may also support many experimental campaigns or product lines. A transparent assessment must state how inputs and burdens are attributed across these functions. Physical allocation, economic allocation, subdivision, system expansion, and substitution-based approaches answer different questions. In ISO-oriented studies, avoiding

DRAFT

Table 37: Evidence required to compare lifecycle burdens of autonomous production policies. A completed calculation in one row does not establish the claims excluded in the rightmost column.

| Accounting interface | Minimum records and units | Comparison requirement | Inference not established alone |
|---|---|---|---|
| Productive function | Conforming units, quality, service life, delivery interval | Identical functional service | Similar output counts imply equal performance |
| Direct operation | Metered energy, water withdrawal/consumption, chemicals, yields | Matched utilization and operating regime | On-site reductions lower total burden |
| Equipment and supply | Hardware bill of materials, upstream processes, replacement frequency | Comparable foreground and background coverage | Supplier location reveals ecosystem impact |
| Recovered flows | Collection, qualified output, treatment losses, destinations | Consistent allocation/substitution scenario | Recycled mass equals avoided extraction |
| Ecological pathways | Pollutant identity, receptor, geography, period, ecological surveys | Valid characterization or field observations | Modeled impact equals ecosystem recovery |
| Authority and review | Evidence provenance, scope decisions, contestation record | Comparable disclosure and institutional rights | Complete records establish legitimate action |

allocation where feasible is a preferred design direction; where it cannot be avoided, the selected treatment and its consequences need to be disclosed and tested [83, 33]. A particular allocation rule can change the reported burden attributed to one product while leaving physical emissions at a facility unchanged.

A distinction between *attributional* and *consequential* questions is therefore required. Attributional assessment asks what share of an existing or modeled network's burdens is associated with supplying the defined service. Consequential assessment asks which physical changes follow from a specified decision relative to an explicit alternative. The latter requires assumptions about displaced processes, marginal suppliers, market responses, production scaling, time horizon, and relevant policies. Neither interpretation should be treated as inherently correct for every research purpose. When comparing an autonomous and a conventional facility, the manuscript proposes reporting matched attributional inventories and, separately where defensible, decision-dependent scenarios. A putative benefit should not be claimed by subtracting a speculative displacement credit from a measured burden without showing the causal conditions under which displacement could occur.

A hypothetical example makes the accounting problem explicit. Suppose two production architectures each supply ten thousand specification-conforming components. Architecture A uses a newly manufactured controller with a longer service life and fewer rejected components; B uses refurbished control equipment but higher rework energy. An assessment that counts only purchased new hardware might favour B, while one that counts all electricity and rework might favour A. Inclusion of upstream process chemicals, equipment lifetime, and qualified end-of-life recovery could further change the comparison. No numbers or rank order for these hypothetical alternatives are asserted here. The scientific task is to obtain actual process data, specify competing plausible substitutions, and test sensitivity to operating duration, utilization, and supply-chain location before reporting a comparative result.

## 23.5 The Location and Timing of Ecological Burdens

Two inventories can have the same aggregated quantity of an elementary flow yet differ in their potential consequences because emissions occur in different watersheds, soils, atmospheric conditions, or seasons. Characterization can be differentiated by location and time where suitable models and measurements exist, as in Equation (58). ISO 14046 defines a water-footprint assessment within an LCA framework, while models such as USEtox address specified toxicological and freshwater ecotoxicological pathways. Neither method can recover missing local ecological observations by formal manipulation alone [85,

DRAFT

186]. A 2024 critical review of 86 spatially explicit LCA studies found persistent constraints involving detailed inventory data and methodological expertise; the review identified a lack of direct comparative evidence that spatialization systematically improves assessment accuracy; regionalization therefore requires an explicit case-specific justification rather than a presumption of superiority [169].

**Proposition 14** (Non-identification of spatially differentiated impact from aggregate flow alone)**.** *Suppose a single elementary flow of total amount $m > 0$ may occur at either of two locations $r_1, r_2$, with distinct applicable linear characterization factors $c_{jkr_1} \neq c_{jkr_2}$ for a specified impact category $j$ and interval. If only $m$ is observed and allocation between locations is unknown, the characterized result for that category is not uniquely identified by the aggregate inventory.*

*Proof.* Consider the two feasible distributions $(I_{r_1}, I_{r_2}) = (m, 0)$ and $(0, m)$. Both have observed total $m$. Under the stipulated linear characterization model, they yield $B_j = c_{jkr_1} m$ and $B'_j = c_{jkr_2} m$, respectively. The results differ because $m > 0$ and the two factors are unequal. Thus aggregate observation of $m$ does not uniquely determine $B_j$. The conclusion is conditional on these available distributions and characterization factors and makes no claim about an observed ecological injury. □

For illustration, an entirely hypothetical model with factors of 2 and 8 impact-units per kilogram would assign respectively 20 and 80 impact-units to ten kilograms of the same specified pollutant, depending on location. The figures are constructed for dimensional illustration; they are neither published fate factors nor observations at real sites. If characterization factors are equal, this particular counterexample does not apply. If effects are nonlinear, location, time, and ecological interactions may be more complicated still. The proposition formalizes a basic information limitation relevant to supplier relocation and environmental burden shifting: aggregate discharge totals cannot by themselves identify spatially differentiated outcomes. Better tracking can narrow the uncertainty only where location-specific data and scientifically adequate impact models are available.

Temporal distributions introduce additional differences. A short thermal discharge during a sensitive reproductive period may not be ecologically equivalent to the same total heat over a longer interval; a transient chemical release can interact with hydrological flows, pre-existing contamination, and recovery lags. A discounted future cost is not a physical measure of delayed biological recovery. Ecological assessments should therefore state temporal sampling resolution, persistence assumptions, latency, exposure duration, and the horizon over which risks are considered. Such choices can determine which species, communities, or future pathways are visible within the study. Section 24 will consider constraints on admissible production trajectories; the present section establishes what must first be counted or measured for those constraints to be empirically interpretable.

## 23.6 Environmental Accounting and Historically Organized Responsibility

The boundaries of economic responsibility are historically constituted. A factory may control the selection of its components and equipment but have limited access to evidence on its suppliers' chemical discharges or mining-related habitat transformations. Ownership, confidentiality, procurement contracts, logistics arrangements, and public disclosure obligations determine which records can be obtained, challenged, corrected, or used in a decision. Historical materialism is useful here because ecological burdens are not just physical quantities transported through technology; they are produced within relations that also regulate the circulation of evidence, costs, authority, and remedial obligations. Different institutional arrangements can expose or conceal different portions of the same material chain.

Accounting rules also have political-economic consequences. Allocation may determine which product category appears more resource-intensive; organizational boundaries can move burdens outside a firm's direct accounts; and a claim of climate efficiency may become a criterion for investment while underreported water impacts remain localized. These possibilities do not establish intentional concealment or imply that one ownership form always performs better. Rather, they motivate comparisons of disclosure quality, auditing, supplier bargaining, and access to remedy across privately owned, public,

DRAFT

cooperative, and commons-oriented arrangements. The same empirical functional unit could be supplied through different institutions, and the resulting lifecycle outcomes should be evaluated independently.

The present GR proposal adds a linked record of production events, access decisions, material pathways, observational provenance, and changes in responsible institutions. Such a record may help reconstruct why a particular supplier, equipment replacement, or waste-treatment route was selected. But established databases, product passports, provenance ontologies, and audited LCA software already capture important parts of this work. GR-specific benefits must be tested against those competent alternatives rather than postulated. One pertinent test asks whether recording institutional restrictions and substitution decisions improves estimates of counterfactual burdens, uncertainty identification, or the ability of affected parties to dispute a claim. A negative finding would limit the justification for a distinct GR layer in environmental accounting.

Ethical and legal judgments remain distinct from inventory arithmetic. Potential climate effects, predicted freshwater deprivation, observed habitat damage, and harms to individual nonhuman organisms involve different kinds of evidence. What constitutes an acceptable intervention or compensable loss requires an explicitly stated normative and institutional process. A global numerical environmental-footprint score does not grant permission for a facility, establish that a harmed community has consented, or prove restoration of ecological generativity. IPBES emphasizes that multiple values of nature can be overlooked when evaluative procedures privilege narrow measurable contributions [60]. The paper therefore reports physical and ecological evidence without presuming a single globally authoritative valuation scale.

## 23.7 Comparative Research Protocol and Identification Limits

A practicable study might compare conventional, optimized, and AI-adaptive production of an otherwise identical precision component across the same service horizon. The protocol would first register the functional unit, output tolerances, utilization, location, equipment service lives, and excluded flows. It would then collect primary process records, supplier-specific data where possible, repair and replacement histories, recovery yields, and logistics pathways. Generic databases would fill declared gaps with uncertainty labels. The comparison should report at least one foreground-only result and one consistently bounded full-lifecycle estimate for each alternative, exposing the contribution of background data instead of hiding it in a headline claim.

A second, spatially differentiated arm would identify one or more impact pathways for which inventory location materially affects exposure, such as water consumption within a stressed basin or emissions of identified substances into receiving waters. A credible ecological claim would require external observations and a model suitable for the relevant receptor. Parallel supplier or time-period comparisons may help, but the researcher would still need to address non-random sourcing, policy changes, background climate, and data quality. A third arm would test the consequences of allocations and counterfactual substitutions, reporting whether any claimed advantage changes sign as assumptions vary. None of the three arms requires that automation outperform its comparator; mixed or null results are scientifically informative.

The decisive methodological caution is that an inventory estimates pressures under declared accounting conventions; a characterization model estimates potential consequences given a model; a field study observes selected ecological outcomes; and an institution decides which interventions may proceed. These stages can be linked, but evidence does not automatically transfer from one to the next. The same caution applies to uncertainty: sampling error, sector aggregation, model-form uncertainty, temporal extrapolation, and the unknowability of some counterfactual actions should be identified rather than combined into a confidence interval without justification. Transparency and sensitivity analysis are essential, yet neither substitutes for measurements of the ecological relations affected by production.

This section therefore supplies the measurement architecture for the next two inquiries. Section 24 examines whether prospective production policies satisfy specified ecological viability constraints under bounded or probabilistic uncertainty. Section 25 studies how output scaling, induced demand, overlap-

DRAFT

Table 38: Proposed sensitivity and falsification programme for lifecycle claims. Indicators remain separate unless a justified decision procedure specifies how to compare them.

| Research question | Required variation or measurement | Candidate indicator | Finding that limits the claim |
|---|---|---|---|
| Does automation reduce lifecycle burden? | Equal service, varying equipment and input histories | Category-specific burden per functional unit | Advantage disappears beyond factory gate |
| Does repair extend productive benefit? | Actual failure, replacement, refurbishing and material-flow logs | Burden under observed service duration | Replacement burden outweighs avoided inputs |
| Does recovery displace primary input? | Supplier evidence, yield, quality, market response | Qualified substitution and residual losses | Collection fails to change primary production |
| Does location change ecological exposure? | Discharge locations, characterized pathways, field samples | Receptor-specific impact and uncertainty | Difference is unidentifiable at data resolution |
| Does GR add useful traceability? | Matched conventional/GR provenance and reviewer tasks | Reproducible decisions and resolved discrepancies | No measurable improvement over good LCA |

ping exposures, and persistent infrastructures change total ecological burdens across space and time. In historical-materialist terms, lifecycle accounting can make the material dependencies and displaced burdens of production more observable; it does not by itself settle how productive relations ought to be organized or which ecological transformations are justified. Its value lies in enabling claims of ecological compatibility to be tested, qualified, and, where warranted, rejected.

# 24 Formal Conditions for Robust Ecological Viability

Section 21 identified the ecological conditions whose continuation must be investigated alongside productive reproduction. Section 22 distinguished observed recovery from the recoverability of ecological capacities, and Section 23 established the physical inventories and impact models through which production-related pressures may be estimated. The present section asks whether a production policy can meet a specified productive service while keeping a coupled technological–ecological system within independently justified constraints under stated disturbances. Its objective is to provide a precise language for feasible, infeasible, and indeterminate cases. These categories are relative to models, observations, horizons, disturbance assumptions, and institutions that authorize the relevant interventions. Mathematical viability can help expose contradictions among production demands and ecological requirements; it neither discovers ethically appropriate thresholds nor certifies the ecological compatibility of an operating facility.

## 24.1 The Scope of a Viability Inquiry

Viability theory concerns trajectories that remain within declared constraints under admissible decisions; it has applications in ecological and economic resource management as well as control engineering [7, 28]. The set-invariance literature examines related questions of constrained state evolution, feedback, and robustness [13]. Béné, Doyen, and Gabay applied viability analysis to a bio-economic system with joint ecological and economic requirements, showing the relevance of compatibility among constraints rather than optimizing a single economic objective [11]. The use of a viability kernel in this paper therefore represents a deliberate application and extension of established mathematical approaches to the production-dependency and ecological-evidence interfaces developed in previous sections. It is not presented as a newly invented theorem or a new universal measure of sustainability.

For an autonomous production installation, the research object is a joint state containing at least the

DRAFT

relevant productive capacity, material and energy stocks, and ecological attributes. A plant can meet its delivery targets while degrading an ecological receptor; conversely, a policy that reduces local exposure may fail to maintain a specified level of productive service. These are distinct failures relative to different constraints. In addition, an actor may possess physically feasible control actions that are unauthorized under applicable law or institutional agreements. The admissible action set must therefore be declared with its physical, operational, and institutionally enforceable restrictions, while recognizing that normative judgments about nonhuman interests and the legitimacy of intervention cannot be inferred from dynamics alone.

The system boundary is an explicit research choice. At a minimum it identifies physical facilities, productive services, supplier and disposal pathways, ecological receptors, affected locations, and the evaluation horizon. The smaller boundary used for machine control may differ from that required for lifecycle accounting. A plant-level model can certify conditional compliance with an on-site temperature constraint but say nothing, by itself, about upstream groundwater withdrawal or chemical discharge. A viability claim concerning the entire productive pathway requires those dependencies to be incorporated through expanded states, bounded external inputs, or clearly qualified supplementary evidence. The model also needs to distinguish robust feasibility for *all modeled disturbances* from performance under the disturbances that happened to be observed.

## 24.2 Coupled State Dynamics and the Meaning of Constraints

Let $t = 0, \dots, H$ index decision intervals over a finite physical time horizon. Denote the coupled state by

$$s_t = (p_t, e_t, h_t) \in \mathcal{S}_t, \tag{60}$$

where $p_t$ contains productive capacities and resource stocks, $e_t$ contains ecological attributes defined at stated receptors, and $h_t$ records any accumulated exposure or delayed-effect variables needed for the model to be sufficiently predictive. The latter is important: a dissolved contaminant may remain ecologically relevant after direct discharge ceases, and a population's response may depend on past recruitment or disturbance. Markovian notation does not remove ecological memory; it requires that the variables needed to represent it be included or that a model with explicit history dependence replace the present formulation. Each component retains its physical units. Production rates, water withdrawals, temperatures, abundances, and habitat attributes cannot be added as if they were commensurate scalar quantities.

A provisional transition model is

$$s_{t+1} = F_t(s_t, u_t, w_t; \theta), \qquad u_t \in \mathcal{U}_t(s_t), \quad w_t \in \mathcal{W}_t(s_t, u_t), \tag{61}$$

where $u_t$ denotes a physically and institutionally admissible production or mitigation decision, $w_t$ a modeled exogenous disturbance, and $\theta$ calibrated or explicitly uncertain physical parameters. For an engineering subsystem, components of $F_t$ may come from measured energy and material balances, qualified equipment capacities, or reliability models. For an ecological subsystem, they may require field-based exposure–response models, hydrological or atmospheric transport models, species-specific demographic information, and evidence on cumulative stressors. The coupled map is hypothetical until it has been identified and tested. Its domain of validity, extrapolation limits, and unresolved causal mechanisms must be recorded. In the recursion below, $\theta$ denotes a fixed, declared parameter value or a parameter vector whose uncertainty has already been included in an appropriately constructed state or disturbance specification. Simply writing an uncertain $\theta$ in the transition equation does not make a disturbance-robust result robust to parameter misspecification.

Let $\mathcal{K}_t \subseteq \mathcal{S}_t$ represent the allowed state configurations at decision interval $t$. It may encode lower limits on productive safety stocks, minimum levels of a defined productive service, or explicit bounds on selected ecological exposures and conditions. Constraints on outputs and actions can instead enter through a state-dependent feasible control set $\mathcal{U}_t^{\mathrm{svc}}(s)$, which includes a specified production target and legitimate operational rules. A terminal set $\mathcal{T}_H \subseteq \mathcal{S}_H$ can require continuation readiness, such as available repair inputs and maintained ecological margins, rather than permitting a policy to exhaust its machinery

DRAFT

or ecological surroundings at time $H$. The admissible state set may vary with season, climate conditions, monitoring results, or formally revised ecological requirements.

The ecological elements of $\mathcal{K}_t$ are scientifically and normatively substantial assumptions. A threshold for a regulated pollutant is not identical to an ecological threshold for a sensitive community; a legally applicable limit is not necessarily sufficient to protect all affected organisms; and a model-based habitat proxy can fail to capture ecological continuation. The literature on Earth-system boundaries illustrates the importance of clearly specified scales and criteria while cautioning against interpreting global thresholds as facility permits [159]. Relevant limits should be chosen through evidence, applicable law, participation, and context-specific evaluation rather than derived from the mere convenience of a mathematical feasible set. As argued in Section 21, a formally satisfiable constraint system may still exclude important ecological interests.

## 24.3 Robust Viability over a Declared Horizon

Under full state observation and a stagewise disturbance model, a feedback policy is a family $\pi = (\pi_0, \dots, \pi_{H-1})$, with $u_t = \pi_t(s_t)$ selected from actions available at time $t$. Decisions are *nonanticipative*: a policy may react to information already available but cannot select an action using a disturbance that has not yet been observed. The finite-horizon robust viability set is defined by

$$\begin{aligned}\mathcal{V}_{0:H}^{\text{rob}} = \{s_0 \in \mathcal{K}_0 : \exists \pi \in \Pi_{\text{ad}},\ \forall (w_0, \dots, w_{H-1}) \in \mathfrak{W}(\pi),\\ s_t^{\pi,w} \in \mathcal{K}_t \text{ for } t = 0, \dots, H, \quad s_H^{\pi,w} \in \mathcal{T}_H\},\end{aligned} \tag{62}$$

where $\mathfrak{W}(\pi)$ contains the disturbance histories allowed by the declared disturbance model when policy $\pi$ is used. Productive-service and authorization conditions are incorporated in $\Pi_{\text{ad}}$ or $\mathcal{U}_t^{\text{svc}}$. The universal quantifier refers only to the disturbances represented in $\mathfrak{W}$: unknown extremes and omitted ecological processes are not made harmless by the definition. In particular, if $\theta$ is an unknown but fixed member of a set $\Theta$, a guarantee over such uncertainty requires an explicit quantifier over $\Theta$ and a policy information structure that forbids selecting an action using its unobserved value. An augmented belief or parameter-state model, or a separately justified robust parameter analysis, may then be necessary; independently reselecting $\theta$ at every stage can be overly conservative and is not equivalent to a fixed unknown parameter. Likewise, if the terminal set lacks a meaningful continuation condition, short-horizon viability may say little about longer-term productive or ecological reproduction.

The familiar predecessor construction becomes available when the state is sufficient to predict the next admissible transitions and the disturbance choices are *rectangular*, meaning the declared set of possible disturbances at each interval can be selected independently of unmodeled cross-time restrictions conditional on the current state and action. Define

$$\text{Pre}_t(A) = \left\{s \in \mathcal{K}_t : \exists u \in \mathcal{U}_t^{\text{svc}}(s)\ \forall w \in \mathcal{W}_t(s,u),\ F_t(s,u,w;\theta) \in A\right\}. \tag{63}$$

The use of a stagewise predecessor is an existing dynamic-programming construction. If uncertainty is coupled across time, hidden persistent parameters or correlated disturbances may need to be introduced into the state or treated through a different information structure. The proposition that follows is therefore a conditional result for the specified fixed-parameter, stagewise rectangular model; it does not separately establish parameter-robust viability. Merely using a computationally convenient independent disturbance envelope can make a test either unduly conservative or incorrectly optimistic.

**Proposition 15** (Finite-horizon viability recursion)**.** *Suppose the fully observed system follows Equation* (61)*, the admissible state and action sets are declared, the disturbance choices obey the stated stagewise rectangular model, admissible feedback selections exist in the required policy class (for example on finite state spaces or under suitable measurable-selection conditions), and feedback actions may be selected after observing the current state but before the current disturbance. Let* $\mathcal{V}_H = \mathcal{K}_H \cap \mathcal{T}_H$ *and define* $\mathcal{V}_t = \text{Pre}_t(\mathcal{V}_{t+1})$ *for* $t = H-1, \dots, 0$*. Then* $\mathcal{V}_0$ *is precisely the set of initial states for which a nonanticipative feedback policy satisfying the specified state, service, and terminal constraints exists against all modeled disturbances over the finite horizon.*

DRAFT

*Proof.* At $t = H$, membership in $\mathcal{V}_H$ is exactly satisfaction of the final state and continuation constraints. Assume that membership in $\mathcal{V}_{t+1}$ is equivalent to the existence of a viable policy for the remaining intervals. If $s_t \in \mathrm{Pre}_t(\mathcal{V}_{t+1})$, an admissible first control exists for which each modeled next state belongs to $\mathcal{V}_{t+1}$; the induction hypothesis supplies a continuation policy after that next state becomes observable. Conversely, if an admissible feedback policy succeeds for all modeled disturbance histories, its first control must place each possible next state in $\mathcal{V}_{t+1}$, so $s_t \in \mathrm{Pre}_t(\mathcal{V}_{t+1})$. Backward induction proves the result. □

The proposition is an elementary finite-horizon result from standard viability and robust-control mathematics [7, 13, 28]. It establishes an equivalence *inside the specified model*, not a theorem that real ecosystems can be controlled robustly or that the modeled constraints fully represent their generativity. Computation of $\mathcal{V}_0$ may be intractable in high-dimensional, hybrid, or poorly observed systems. A constructed observability counterexample makes the assumption consequential. Suppose two ecological states share the same measured temperature and discharge record but differ in an unobserved cumulative contaminant stock. In the first state, a policy that maintains production respects the registered threshold; in the second, the same policy crosses a delayed ecological limit, while an immediate reduction is required. A feedback law computed with the full contaminant stock can distinguish these cases, whereas a controller restricted to the shared observation cannot select different actions on that basis. If the admissible actions needed to protect the two states are mutually exclusive, robust feasibility under the restricted information structure may be empty even though each fully observed state has an individually viable action. This is a hypothetical illustration of information dependence; it does not establish that the described conflict occurs at a specific site. Approximations can change the direction of the guarantee: an unverified outer approximation may label infeasible states viable, whereas a certified inner approximation can support a narrower conditional sufficiency claim.

## 24.4 Incompatibility between Productive Service and Ecological Protection

The viability construction permits empty feasible sets. Consider a deliberately simplified dimensionless ecological state $e_t$, representing a single habitat-condition indicator, with dynamics

$$e_{t+1} = e_t + 0.4 - u_t - w_t, \qquad e_0 = 5, \quad u_t \in [0.5, 1], \quad w_t \in [0, 0.2]. \tag{64}$$

Here $u_t \geq 0.5$ is a hypothetical minimum production intensity required to deliver a specified service, $0.4$ is a hypothetical background regenerative increment per interval, and $w_t$ is an independently selected adverse exposure in a declared uncertainty range. Let the ecological constraint be $e_t \geq 4$ throughout $H = 4$ intervals. No numerical quantity in this example is derived from a field observation, industrial process, or published ecological threshold; its purpose is to make the logic of conflicting constraints transparent.

Because the disturbance sequence $w_t = 0.2$ is admissible and the smallest permissible production intensity is $u_t = 0.5$, every admissible policy obeys

$$e_4 \leq 5 + 4(0.4 - 0.5 - 0.2) = 3.8 < 4. \tag{65}$$

Thus the specified initial state lies outside the robust viability set: production requirements and ecological conditions cannot jointly be satisfied over four intervals against the declared disturbance sequence. This does not establish impossibility of production in a real ecosystem. The finding changes if the production requirement is reduced, the system can substitute less damaging processes, ecological responses are estimated differently, mitigation modifies the dynamics, or the disturbance assumptions are revised on evidence. Indeed, with otherwise identical hypothetical assumptions, a two-interval horizon leaves a nonempty margin under the least intensive policy while the four-interval problem is infeasible. The difference illustrates why selecting a convenient horizon can conceal a longer-term incompatibility.

This is an example of a constrained decision problem, not an optimization failure that can always be repaired by tuning weights. If no control satisfies non-negotiable physical or ecological conditions, no

DRAFT

Table 39: Illustrative claims in an autonomous-production viability analysis and evidence required beyond mathematical calculation. Each row concerns a different interface.

| Claim examined | Mathematical or empirical object | Evidence required | Unsupported inference to avoid |
|---|---|---|---|
| Productive feasibility | State dynamics, qualified output and input stocks | Process logs, supplier and repair capacities | A simulated target proves operational autonomy |
| Ecological state compliance | $e_t \in \mathcal{K}_t$ across modeled trajectories | Local exposure data and relevant receptors | A numerical limit proves ecological wellbeing |
| Robustness | Disturbance envelope $\mathcal{W}_t$ and error bounds | Historical extremes, sensitivities, independent stress tests | Bounded-model success covers unknown extremes |
| Monitoring adequacy | Observation map and unobserved state set | Sensor validation, detection limits and field coverage | A measured proxy identifies whole ecosystem |
| Institutional admissibility | Authorized actions and declared decisions | Applicable law, access rights, participation records | Technical feasibility supplies permission |
| Continuation readiness | Terminal $\mathcal{T}_H$ and longer-horizon tests | Ecological recovery and infrastructure replacement evidence | Short-horizon survival proves indefinite continuation |

scalar objective weighting can make an infeasible action admissible. A researcher may investigate alternative production volumes, locations, technologies, recovery activities, or institutional arrangements; each such change defines a new scenario requiring its own physical and ecological accounting. Claims of impossibility must therefore remain conditional on the model and what interventions were legitimately included. Conversely, a feasible numerical trajectory is only a model-based existence claim until the dynamics, observations, and constraints have been validated.

## 24.5 Partial Observation and Uncertain Ecological Models

The full-state recursion is an analytic benchmark, not a realistic description of most ecological management. Autonomous plants generally observe some physical process variables with high temporal resolution, while ecological attributes may be measured irregularly, indirectly, or at spatial scales different from relevant habitat processes. Denote observations by $y_t = O_t(s_t, \nu_t)$, with sensor or sampling uncertainty $\nu_t$. A feasible implementable policy has the form $u_t = \mu_t(y_0, \ldots, y_t, u_0, \ldots, u_{t-1})$, subject to information-access and legal restrictions. A state-feedback policy that requires inaccessible ecological states cannot be implemented simply because a mathematical optimizer can compute it.

One conservative method is to maintain a set $\mathcal{I}_t$ of states consistent with available observations and declared uncertainty. An action may be called robustly admissible for the information state if it meets its constraints for every state in $\mathcal{I}_t$ and every modeled disturbance; the information set must then be updated after each observation. A Bayesian alternative uses a declared belief distribution and probabilistic criteria, but its conclusions depend on prior, likelihood, and model adequacy. Neither uncertainty representation is universally appropriate. Sparse ecological sampling, observation-induced disturbance, nonrandom access to sites, and unrecognized species interactions can invalidate apparently precise estimates. In some cases the defensible result is an unresolved feasibility judgment rather than a numerical certificate.

Chance constraints answer another question. A condition of the form

$$\Pr_{\mathcal{M}}^{\pi}\left(\bigcap_{t=0}^{H}\{s_t \in \mathcal{K}_t\}\right) \geq 1 - \alpha \tag{66}$$

asserts a joint-horizon probability bound under a specified distributional model $\mathcal{M}$ and explicitly fixed admissible policy $\pi$; neither the disturbance distribution nor the policy is supplied by the inequality itself.

DRAFT

It does not imply that every admissible disturbance trajectory is safe. Nor does a family of marginal probability bounds on individual constraints automatically provide the same joint guarantee; dependence and allocation of the total tolerated risk must be addressed. What constitutes an acceptable $\alpha$ is an institutional and ethical question that cannot be settled by the formalism. When relevant ecological risks concern irreversible losses or the interests of organisms inadequately represented in the modeled state, a small simulated violation probability may be an insufficient basis for authorization.

Model-form uncertainty deserves separate treatment from stochastic noise. A computer may predict discharge temperature precisely while omitting a relevant downstream habitat, or may predict water availability while failing to represent supplier withdrawals. Sensitivity analysis across alternative ecological mechanisms, spatial boundaries, recovery times, and intervention histories is consequently essential. Demonstrated agreement within historical training data does not by itself validate extrapolation to altered climate or a new production architecture. If models disagree about whether a critical ecological constraint is satisfied, that disagreement should remain visible in the reported viability classification.

## 24.6 Productive Reproduction and the Horizon of Ecological Continuation

Robust viability over a finite horizon is different from the capacity for continued reproduction. A control policy could keep an ecosystem indicator above a threshold through one production cycle while exhausting qualified spare parts, reducing regenerative capacity, or accumulating unobserved stressors. A terminal readiness set can require particular production stocks and ecological margins at $H$; yet no terminal condition can guarantee indefinite viability without conditions on post-horizon dynamics and feasible future interventions. An infinite-horizon controlled-invariant set is a distinct object, requiring stronger and often empirically inaccessible assumptions about long-term disturbance and ecological dynamics [13, 7].

The relationship with Section 22 is especially important. A control policy may remain within a declared threshold during ordinary operation but still undermine the system's recoverability following an accident. Conversely, a restoration programme may increase the range of recoverable conditions while imposing short-run material and energetic costs. Ecological viability and recoverability should therefore be compared under the same spatial units and model histories but reported as distinct properties. It would be misleading to present viability of a selected control variable as equivalent to the persistence of ecological assemblages, the wellbeing of particular nonhuman organisms, or the preservation of open-ended generative possibilities. Ecological futures can depart from the modeled state space, and their significance may resist a fixed terminal-score representation.

The historical-materialist dimension concerns who can alter these possibilities. Production quotas, access to alternative suppliers, rights to stop operations, environmental monitoring budgets, and responsibilities for restoration arise within historically constituted production relations. Such arrangements change both the set of feasible controls and whose evidence enters the decision process. A publicly owned, investor-owned, cooperative, or commons-oriented facility may confront different cost constraints and duties, but none is guaranteed ecological viability by its ownership form. The model makes those differences investigable by identifying decision rights, productive capacities, and environmental constraints separately. It does not reduce relations of production to economic coefficients or presume that a single global evaluator can resolve heterogeneous claims concerning ecological life.

## 24.7 Comparative Research Design and Failure Criteria

An empirical study could compare several control and production architectures supplying the same quality-qualified industrial service over a registered horizon. A baseline facility, a conventional energy- and material-optimized policy, and an AI-assisted adaptive policy would use matched physical and institutional boundaries wherever possible. The evaluation would first document actual electricity, water, material, maintenance, and production records. It would then independently estimate relevant ecological states and exposures using field measurements and justified mechanistic models. Only after such

DRAFT

Table 40: Proposed tests capable of limiting or rejecting claims about robust ecological viability in autonomous production.

| Research question | Comparison or validation design | Reporting unit | Finding that limits the claim |
|---|---|---|---|
| Can service and ecology coexist? | Joint production and receptor observations across the same horizon | Qualified service and bounded ecological indicators | Feasible set empty for justified assumptions |
| Does robust control outperform ordinary control? | Matched services, uncertainty tests and out-of-sample disturbances | Constraint frequency, output, failures and uncertainty | No performance gain over competent baseline |
| Are ecological proxies adequate? | Independent receptor sampling, alternative observation maps | Prediction error, detection rate and missing states | Unobserved deterioration invalidates compliance proxy |
| Does a longer horizon change feasibility? | Recompute viable sets for multiple horizons and terminal conditions | Horizon-dependent feasible-state set | Apparent advantage disappears with repair or recovery needs |
| Can GR improve auditability? | Compare existing LCA inventories and operational-technology provenance with additional relational records | Reproducibility of decisions and resolved disputes | No measurable improvement over established methods |

measurement would a viability algorithm operate on the bounded dynamics. A policy should be classified as robustly viable only within the declared model, uncertainty envelope, thresholds, and observation assumptions; broader empirical ecological compatibility would require additional independent evidence.

A stronger protocol would preregister productive functional units, ecological receptor selection, uncertainty envelopes, constraint provenance, failed and successful model versions, and any human overrides. Observed ecological change must be compared with plausible alternative causes, including background climate variation, other industries, seasonal cycles, and pre-existing disturbance. The analyst should report the consequences of changing modeling boundaries, disturbance support, thresholds, terminal conditions, and feasible production demand. Since such changes can reverse a feasibility conclusion, model sensitivity is a central result rather than an appendix to a binary viability label.

The GR research proposal would add dated records linking productive events, material substitutions, ecological measurements, model revisions, contested constraints, and responsible institutions. Its potential contribution is the traceability of changing relations and the identification of which choices altered feasible continuations. That contribution is only persuasive where matched evaluation demonstrates an advantage over competent industrial-control logs, provenance standards, and lifecycle systems. Such evidence would also help distinguish an ecological violation caused by an unforeseen disturbance from one caused by a chosen production policy or omitted measurement pathway, without attributing intention or responsibility solely from a mathematical optimization result.

This section thus closes the first formal bridge between the production-dependency architecture of Sections 6–20 and the ecological conditions examined in Sections 21–23. Robust viability is a disciplined method for asking whether jointly specified productive and ecological constraints can be satisfied across uncertain trajectories. Its scientific use depends on accurate physical models, independently warranted ecological variables, honest information limits, and legitimate institutional procedures. A nonempty calculated viable set neither proves ecological justice nor establishes perpetual reproductive capacity; an empty set instead identifies a conditional conflict whose causes and available alternatives deserve investigation. Section 25 now turns to another consequential limit of facility-level feasibility: autonomous production may scale, relocate, or induce additional material activity, altering cumulative ecological burdens beyond the tested boundary and horizon.

DRAFT

# 25 The Cumulative Effects of Autonomous Production at Scale

Sections 21–24 investigated ecological continuation, recoverability, lifecycle burdens, and the conditional feasibility of production within stated ecological constraints. Each inquiry depends on its system boundary and the level of productive activity assumed. The present section investigates how those conclusions may change when autonomous production becomes less costly, more widely adopted, geographically distributed, or more intensive over time. Its purpose is to distinguish improvements within a particular process from changes in the total biophysical demands and ecological exposures generated by a historically expanding system of production. The analysis draws upon the established literature on rebound effects, industrial material demand, cumulative-impact assessment, and demand-side mitigation. The proposed contribution is an explicit connection among changing productive dependencies, institutional decisions regarding scale, and the evidence needed to judge ecological continuation across multiple interacting sites. Greater autonomy is neither assumed to cause resource growth nor treated as sufficient for ecological improvement.

## 25.1 Distinguishing Productive Efficiency from Absolute Ecological Change

A production cell may use fewer kilowatt-hours, litres of water, or kilograms of material per conforming product while the enterprise, sector, or regional industrial system draws more resources in total. This elementary possibility matters especially where autonomous equipment operates for longer intervals, permits lower unit costs, enables additional product varieties, or is replicated across facilities. The relevant outcome depends on what service is supplied, which environmental flow is counted, and how the observational horizon is drawn. Even the meaning of a "unit" requires care: the same number of items may provide different service lives, reliability, or downstream functions. An efficiency comparison must therefore use a defensible functional unit and acknowledge changes in output quality and service performance.

For a single physically homogeneous flow, define $Q_t > 0$ as the number of functionally comparable units supplied in interval $t$, $i_t \geq 0$ as the flow requirement per unit, and $C_t \geq 0$ as additional flow attributable to capacity expansion, equipment replacement, and other infrastructure processes not already included in $i_t$. In units matched to the selected flow, total requirements within a declared boundary are

$$M_t = Q_t i_t + C_t. \tag{67}$$

The identity is an accounting decomposition, not an empirically estimated causal response. $C_t$ can represent additional electricity or mass only if the selected inventory category and units are consistent; energy, greenhouse-gas emissions, and water consumption require distinct ledgers. If embodied equipment requirements are already amortized in $i_t$, they must not be added again in $C_t$. Where there are many products and flows, the corresponding representation is a typed vector or matrix of inventories, with separate functional units and allocation rules as in Section 23. It is not legitimate to add heterogeneous ecological harms to obtain a single physical total without a scientifically and normatively defended characterization.

A simple constructed comparison illustrates the problem. A factory initially manufactures 100 equally specified units at 10 kWh per unit, with no additional infrastructure electricity included in the accounting boundary; its counted requirement is 1,000 kWh. An alternative automation strategy reduces operating electricity intensity to 7.5 kWh per unit but expands output to 150 units. At the same narrow operating boundary, the counted requirement rises to 1,125 kWh, an increase of 12.5%, notwithstanding a 25% intensity improvement. These numbers are hypothetical and do not demonstrate that automation induced the output expansion. The result could also change after accounting for equipment manufacture, process heat, quality losses, regional electricity mix, or the services delivered by the additional units. The purpose of the example is to show what evidence must be collected before presenting technical efficiency as an absolute ecological gain.

**Proposition 16** (Threshold for an absolute reduction under fixed ancillary requirements)**.** *Let* $Q_0, Q_1, i_0, i_1$ *be strictly positive and let* $C_0 = C_1 = C \geq 0$ *represent unchanged ancillary re-*

*quirements measured in the same physical unit. If $i_1 < i_0$, then total requirements decrease from interval $0$ to interval $1$ if and only if*

$$\frac{Q_1}{Q_0} < \frac{i_0}{i_1}. \tag{68}$$

*Equality produces unchanged total requirements; the reverse strict inequality produces an increase.*

*Proof.* Under the stated accounting boundary, $M_1 - M_0 = Q_1 i_1 - Q_0 i_0$ because the ancillary terms cancel. Since $Q_0 i_1 > 0$, the inequality $M_1 < M_0$ is equivalent to $Q_1/Q_0 < i_0/i_1$. Equality and the reverse case follow from the same identity. The statement makes no causal claim about demand, prices, or productive policies. □

The proposition is an elementary product identity. Its importance here is methodological: an empirical study needs observations of both production intensity and comparable service volume to establish whether a particular flow has decreased. When $C_1 \neq C_0$, an additional term $C_1 - C_0$ enters the comparison and the displayed threshold is insufficient. The result therefore cannot establish lifecycle ecological compatibility and is offered neither as a new economic law nor as a substitute for the viability conditions in Section 24.

## 25.2 How Efficiency Changes Can Alter Productive Demand

The rebound-effect literature examines circumstances in which improvements in technical efficiency generate behavioural or economic adjustments that offset some anticipated resource savings. Historical discussion commonly traces the problem to Jevons's analysis of coal consumption, while contemporary studies distinguish direct responses, indirect expenditure or input substitution, and economy-wide adjustments. The present paper adopts these distinctions as an existing analytical vocabulary and does not infer rebound merely because an absolute aggregate rises. Such inference requires a counterfactual showing how much resource use would have occurred without the intervention, with comparable technology, output, prices, and relevant exogenous conditions.

For a given resource and consistent time boundary, a conventional rebound metric can be written as

$$\mathcal{R} = \frac{S^{\text{eng}} - S^{\text{real}}}{S^{\text{eng}}}, \qquad S^{\text{eng}} > 0, \tag{69}$$

where $S^{\text{eng}}$ is an explicitly modeled engineering saving under fixed specified demand, and $S^{\text{real}}$ is the causally attributable saving against a suitable no-intervention counterfactual. The terms require a common accounting boundary; the engineering saving is not an automatically observable benchmark. A positive $\mathcal{R}$ indicates partial offset under that convention, while $\mathcal{R} > 1$ denotes backfire. Negative values are also possible if accompanying responses create additional savings. These categories must be estimated, not read from the mere coexistence of increased automation and increased resource use.

In AI-mediated production, potential channels include direct expansion of output when lower production costs increase demand, reallocation of savings to other energy- or material-intensive processes, capital investment in additional lines or computing infrastructure, faster replacement of product generations, and changed supply-chain organization. There can equally be countervailing effects: automated defect detection may reduce rejected products, improved scheduling may reduce peak demand, qualified remanufacturing may extend equipment life, and production planning may limit unnecessary inputs. The sign and magnitude of net change therefore remain empirical questions. In addition, a firm's output growth might have occurred under either control architecture; attributing this growth to automation without a comparison design would overstate its causal effect.

Evidence from other sectors supports taking rebound seriously without assuming a transferable universal percentage. Schütt and colleagues' meta-regression of 43 studies and 1,118 estimates reported heterogeneous household-level rebound effects. Under the authors' standardized reference case of an energy-efficiency improvement in residential energy use, specifications yielded mean total microeconomic rebound estimates of approximately 41–52% [166]. That is household evidence and should not

DRAFT

Table 41: Distinct mechanisms through which scaling autonomous production may change total burdens. The entries identify tests and possible outcomes, not universal effects.

| Mechanism | Required observations | Appropriate comparison | Inference to avoid |
|---|---|---|---|
| Unit-intensity change | Energy/material use per qualified service; utilization, yield | Matched process and output specification | Lower intensity establishes lower totals |
| Output or service growth | Volumes, prices, product quality, demand shifts | Counterfactual service trajectory | Co-occurring growth proves rebound |
| Capital-stock expansion | Equipment, buildings, chips, lifespan, embodied inventories | Amortized and total buildout burdens | Operational-only inventory is complete |
| Supply relocation | Supplier-specific flows, transport, technologies, recipients | Full affected production network | Local decline equals global decline |
| Ecological accumulation | Exposure history, receptor status, stressor interactions | Repeated measurements and model validation | Equal totals yield equal ecological effects |
| Institutional mediation | Investment decisions, access rights, monitoring, thresholds | Matched technical alternatives across governance forms | Ownership category alone determines outcome |

be inserted as a parameter for autonomous factories. Lu and Wang's 2024 firm-level study reported direct energy rebound estimates of 33.5% for manufacturing and 57.9% for services within its studied Chinese samples and methods [115]. Those results illustrate the relevance of production-side evaluation but cannot establish equivalent effects for robotic production, another region, or a different energy policy. Jarvis and King (2024), in a distinct economy-wide modelling study, fit a nonlinear energy–output framework to historical global aggregate series; that model's rebound estimates depend on its economic and physical assumptions and are not a meta-analytic average or an experimentally identified effect of automation [105]. These macroeconomic modelling results and Schütt and colleagues' household meta-regression constitute different forms of evidence; neither licenses transferring a fixed rebound coefficient to autonomous production.

Three causal questions should be kept separate: whether autonomy changes per-unit requirements; whether autonomy causes a change in the volume or composition of final productive service; and whether the resulting changes alter physical and ecological burdens across the full affected network. The first can often be assessed with matched process measurements. The second requires price, investment, demand, and institutional evidence, including relevant counterfactuals. The third additionally requires upstream inventories, changing energy or material supply, and ecological receptor information. The separation prevents a narrow measurement of energy efficiency from being converted into an unsupported claim about economy-wide or ecological results.

## 25.3 The Expansion of Productive Infrastructures

Large-scale adoption involves more than replicating software policies across machines. Additional production capacity can require robot cells, grid connections, transformers, cleanrooms, machine tools, computing facilities, water infrastructure, roads, logistics assets, and qualified service networks. Such capacity is a form of historically accumulated fixed infrastructure: its initial construction concentrates material inputs, while its use and replacement generate further flows over time. Efficiency per completed industrial task may improve even as the material stock needed to support a growing number of tasks expands. Section 21 established why the ecological consequences of these flows cannot be inferred solely from plant uptime; the present question is how many facilities, how much supporting capacity, and which new dependencies are reproduced at scale.

Contemporary indicators show the relevance of this question without measuring a completed transition to self-reproducing AI industry. The IEA's 2026 assessment estimates global data-centre electricity demand at 485 TWh in 2025 and projects around 950 TWh in 2030 under its central assumptions [68]. These figures concern *all* data-centre electricity use, not exclusively factories employing autonomous

DRAFT

AI. They indicate that an enabling infrastructure may grow even as performance per computational task changes. Likewise, UNEP's *Global Resources Outlook 2024* documents the historical rise of global material extraction and analyzes prospective resource use under alternative policy pathways [184]. Neither report isolates the contribution of autonomous production to resource demand. The reports instead establish why prospective assessments must account for infrastructure growth and indirect resources rather than extrapolate from component-level efficiency.

The IPCC's Sixth Assessment Report emphasizes that industrial mitigation depends upon a combination of energy and material efficiency, circular material flows, altered demand, and technological changes in primary production [58]. Its analysis cautions against relying solely on incremental process improvements where material demand continues to grow. For autonomous industry, this suggests treating productive expansion as an explicit explanatory variable alongside energy mix, manufacturing yields, equipment lifetime, and recovery. Even if computing electricity demand grew exactly as projected, attributing the growth to an industrial automation policy would require finer sector-specific evidence; similarly, an aggregate material-throughput trend does not by itself identify the induced-demand effects of robots.

Material stocks also introduce delays. Newly installed equipment may remain in service for years, but replacement intervals, repairability, supply-chain qualification, and decommissioning determine its future demands. A policy that maximizes utilization may extend effective service per machine in one setting or accelerate wear and replacement in another. The relevant comparison is cumulative service delivered over a stated horizon with observed maintenance and replacement, not merely annual purchases or installed capacity. Physical additions cannot simply be subtracted using a monetary salvage value; qualified recovery should be assessed through the stock and material-flow conditions in Section 20. Sunk infrastructure may further shape investment and operational choices, making historically constructed production trajectories path-dependent without making their future evolution predetermined.

## 25.4 Cumulative Exposure across Locations and Time

Ecological cumulative effects require information about overlapping pressures, affected receptors, changing background conditions, and the temporal persistence of past interventions. A single facility's environmental inventory may omit the joint influence of nearby operations, upstream mines, electricity plants, transport corridors, or historic discharges. The International Finance Corporation's cumulative impact assessment handbook offers a six-step framework for investigating such overlapping effects and related management responsibilities [72]. It is a practical assessment methodology, not a universal formula that converts industrial inventory totals into biological outcomes.

To make the evidential problem explicit, let $I_{j,r,t}$ be a dimensioned intensity or amount of stressor $j$ released at location $r$ in interval $t$. For a receiving location $a$, a provisional linear transport-and-persistence representation is

$$X_{j,a,t} = X^{\text{base}}_{j,a,t} + \sum_{r} \sum_{\tau=0}^{t} K_{j,a\leftarrow r}(t,\tau)\, I_{j,r,\tau}, \tag{70}$$

where $K_{j,a\leftarrow r}$ is a stated, physically calibrated transfer or persistence coefficient, with units chosen so $X_{j,a,t}$ has an interpretable receptor-specific unit. The equation is an illustrative linear approximation, not a claimed universal ecological law. It may be inappropriate for reactive chemicals, biological accumulation, threshold responses, multiple interacting contaminants, or ecosystems with strongly nonlinear feedbacks. In such cases, a mechanistic nonlinear model or direct field-based comparison is needed. A single value of $X$ is not equivalent to a finding of reduced abundance, impaired reproduction, or lost ecological generativity.

The record of exposure may need to include multiple stressors and their histories. Hydrological extraction, temperature changes, noise, light, airborne emissions, and habitat alteration cannot necessarily be aggregated additively into a scalar risk index. A receptor may be simultaneously exposed to pressures from multiple independently permitted facilities; baseline trends, climate variability, species recruitment, and restoration efforts may affect its response. Seasonal conditions matter where reproductive periods

DRAFT

or migration make otherwise similar exposures biologically nonequivalent. Studies with only annual aggregate loads or short observations can fail to identify these mechanisms even when their inventory arithmetic is correct.

Such effects can also be spatially displaced. Increasingly autonomous fabrication in one region may require refined materials, electricity-generation equipment, and waste processing in distant locations. The physical inventory methods of Section 23 can help track these flows, but a cumulative receptor assessment must additionally characterize transport, persistence, ecological sensitivity, and prior exposures. A local regulatory threshold cannot automatically be interpreted as a sufficiently protective cumulative regional threshold. Where data are sparse, uncertainty should be shown rather than hidden inside a supposedly precise combined ecological score.

## 25.5 Historically Constituted Decisions about Productive Scale

The expansion of productive scale is organized through historically specific relations of investment, ownership, labour, access to technologies, and public authority. A firm may adopt autonomous equipment to lower cost or meet demand; a public enterprise may expand capacity under an infrastructure mandate; a cooperative may emphasize continuity of member livelihoods; and a commons-oriented installation may organize access and maintenance through negotiated obligations. These institutional types do not determine the ecological consequences by themselves. Their significance lies in the incentives and effective decision rights through which feasible production and investment trajectories become realized. The same engineering efficiency gain can therefore coexist with different output decisions, replacement schedules, disclosure practices, and safeguards.

Historical materialism helps connect the microeconomics of scale to the reproduction of production relations. A reduction in direct labour requirements may be associated with investment in additional machinery, a different organization of maintenance work, or new dependence on a limited supplier network. Accumulated productive knowledge can enter machinery that increases output while changing who can command that output. Where costs of extraction, pollution, or habitat loss are not represented in financial decisions or are allocated to other groups, private expenditure and ecological burden may evolve differently. Such divergences require specific causal and institutional investigation; they cannot be inferred from the mere presence of capitalism, public ownership, or autonomous software.

A proposed GR inquiry would represent these relationships as evolving productive possibilities, including which actors may authorize expansion, restrict production, substitute inputs, revise ecological limits, or contest reported outcomes. This can extend the physical accounting developed in Section 23 by preserving the history of decisions that changed scale, location, and supplier access. Yet existing environmental accounts, enterprise systems, lifecycle databases, and cumulative-impact registers already record many of these facts. The additional representation is justified only if it measurably improves causal reconstruction, sensitivity analysis, comparison of feasible alternatives, or contestability relative to competent established practices.

The normative issue is equally significant. A choice to increase industrial output is not automatically justified because a model finds it feasible under selected ecological constraints. The source and legitimacy of those constraints, the representation of affected communities and nonhuman interests, and the possibility of challenging adverse decisions remain separate questions. A model may reveal incompatibility between production expansion and declared ecological conditions. Resolving that incompatibility requires decisions about productive priorities, mitigation, alternatives, and limits, not merely additional computation.

## 25.6 Designing Comparative Tests of Scale-Dependent Burdens

An empirical study should distinguish at least four levels: (i) a matched production operation, (ii) a complete facility and its infrastructure, (iii) the affected supply and consumption network, and (iv) ecological receptors with relevant exposure histories. For each level, the study must specify the service denominator, observation period, system boundaries, permitted substitutions, measured and modeled impacts,

DRAFT

Table 42: A prospective evidence architecture for comparing scale-dependent effects of autonomous production. Each comparison requires an explicit counterfactual and reports unresolved uncertainty.

| Question | Measured or documented variables | Method or baseline | Potential contradictory result |
|---|---|---|---|
| Do unit savings persist? | Service-adjusted input use, scrap, uptime, repair | Matched qualified lines | No real intensity improvement |
| Does output respond causally? | Quantity, pricing, product mix, demand, adoption timing | Difference-in-differences or justified counterfactual | Output shifts explained by non-AI trends |
| Do absolute burdens fall? | Operational and embodied flows, equipment growth, end-of-life | Whole-system material/energy accounts | Aggregate burden rises despite local saving |
| Are burdens transferred? | Supplier flows, region, transport, electricity mix | Multiregional/physical tracing | Local gains remain gains after expansion |
| Do ecosystems continue? | Exposure history, biological indicators, reference observations | Receptor-specific cumulative assessment | Predicted ecological deterioration is absent |
| Does GR tracing add value? | Error rates, missing causal links, dispute resolution | Conventional provenance and LCA baseline | No measurable gain from additional model |

and missing records. An engineering before-and-after comparison can estimate a direct intensity change but generally cannot identify market-mediated rebound. A facility-level energy series can establish net measured consumption but not isolate the cause of its changes. A multiregional input–output model can approximate displaced upstream flows but cannot by itself demonstrate a species-level ecological response.

Where ethically and operationally appropriate, production facilities may be compared during staggered automation adoption, with pre-intervention trends, relevant control installations, measured final service, changing input prices, and output quality taken into account. Alternative designs include within-plant comparisons of qualified lines, natural experiments induced by deployment timing, and structurally calibrated counterfactual scenarios. Site selection must consider that plants implementing more advanced automation may already differ in management, profitability, energy contracts, or product markets. Interference is likely: automation at one firm may shift orders, prices, or emissions to another. Thus the conventional assumption that each plant's treatment leaves its comparators unaffected must be defended rather than presumed. Distributional and institutional data should identify who experiences the resource savings, added production, labour reorganization, or ecological pressure.

The empirical programme should report both favourable and unfavourable outcomes. A reliable finding that autonomous planning reduces absolute resource use after accounting for output responses and infrastructure expansion would count against any general presumption of a rebound penalty. Demonstrating that local exposure declines without higher upstream burdens would weaken a proposed relocation mechanism. Conversely, repeated evidence that additional throughput more than offsets efficiency gains would require stronger attention to productive scale in the viability analysis. Where credible ecological outcome data are unavailable, the study should report bounded exposure estimates or unresolved risk rather than declare the intervention compatible or incompatible with ecological continuation.

The identification of a cumulative effect is also distinct from allocating responsibility for that effect among interacting actors. Physical contributions, legal duties, economic control, and ethical accountability may differ. Comparative research should disclose attribution assumptions and give affected parties access to the information needed to contest them. When cumulative risks depend upon multiple independent enterprises, the appropriate unit for monitoring and prospective coordination may be an ecological region or a shared infrastructure network rather than an individual firm.

DRAFT

### 25.7 Implications for the Reproduction of Productive Conditions

The section establishes that the ecological appraisal of autonomous production must examine absolute flows, productive expansion, geographically distributed infrastructures, and historically accumulated exposure alongside task-level efficiency. Its mathematical results are accounting identities and conditional implications whose scientific use depends on defensible measurements and counterfactuals. The evidence on industrial rebound demonstrates the importance of investigating behavioural and production responses but provides no universal effect size for autonomous AI industry. Cumulative-impact approaches help specify the relevant ecological receptors without guaranteeing that available data can identify all interacting effects.

The historical-materialist problem is consequently concerned with the organization of scale as well as the technical capacity to produce. Who directs investment, whose needs are served by additional output, which inputs become critical, and who bears cumulative ecological burdens may change as productive autonomy advances. The proposed GR representation offers one possible way to link those decisions with evolving physical dependencies, but its added value remains a comparative empirical question. Section 26 will investigate how different institutional forms shape feasible productive trajectories; the subsequent methodological sections will examine counterfactual comparison, case measurement, and the limits of inference. Those inquiries are necessary before the paper's historical and ecological hypotheses can be interpreted as claims about particular real-world production systems.

## 26 Institutional Control over Feasible Production

The preceding section established that technical efficiency and ecological compatibility depend partly on decisions about the scale, location, and duration of production. Those decisions are exercised through historically constituted institutions rather than inferred from an engineering feasibility calculation. This section investigates how ownership, financing, access rights, effective decision-making authority, and accountability influence the productive trajectories that can be implemented with comparable technical resources. Its analytical purpose is to connect the material limits developed in Sections 10–25 to the organization of productive power, while making institutional differences suitable for comparative empirical investigation. Investor-owned, publicly operated, cooperative, and commons-governed arrangements are considered as variable institutional configurations, not as exhaustive social types or a predetermined ranking of ecological outcomes. The central distinction is between a trajectory that is physically possible within a declared model and one whose realization is authorized, financed, coordinated, and sustained by the relevant social relations.

### 26.1 Locating Control within the Reproduction of Production

Historical materialism examines productive forces together with the social relations through which they are reproduced. Machinery incorporates the knowledge and work of different groups, yet its deployment and the disposition of its outputs may be controlled through property, employment, credit, procurement, or public authority. The same industrial robot, materials-handling system, trained controller, and maintenance capability can accordingly operate within distinct relations of production. The differences become analytically consequential when an operator chooses between maintaining a component and replacing it, supplying an essential service and increasing a profitable product line, disclosing a material dependency and keeping it proprietary, or reducing throughput after an ecological threshold is approached. Each choice requires investigation of both physical alternatives and the authority to choose among them.

The literature already supplies sophisticated accounts of these questions. Grossman and Hart distinguish specified contractual rights from residual control when contracts are incomplete, and Aghion and Tirole distinguish formal decision rights from the effective authority created by unequal information [51, 4]. The resource-dependence perspective examines organizational constraints imposed by actors controlling indispensable external resources, while global value-chain scholarship differentiates forms

DRAFT

of interfirm coordination and supplier authority [153, 47]. These approaches precede and partly overlap the proposed generative-relational (GR) inquiry. Its prospective contribution is a typed, temporally inspectable connection between changes in technical dependencies and changes in effective control. That contribution remains a hypothesis about the usefulness of representation, rather than a demonstrated theoretical superiority over established institutional economics, production-network analysis, or historical research.

For an empirical inquiry, control should be disaggregated rather than assigned to a single nominal owner. Relevant dimensions include: title and residual rights over machines and facilities; authority to contract for inputs and electricity; access to process knowledge, models, and operational data; permissions to modify the controller or override an automated decision; authority over maintenance, shutdown, and reinvestment; and the ability of workers, public regulators, suppliers, and affected communities to obtain information and seek correction. A lender may have no day-to-day operating role yet constrain replacement investment through debt covenants. A software vendor may own neither the factory nor its products yet determine when a licensed controller can be updated. Public ownership may coexist with outsourced procurement, proprietary software, and managerial discretion. Member ownership may coexist with severe credit constraints or dependence on a single buyer. These are analytically different channels and require different evidence.

## 26.2 Financing Conditions and the Selection of Productive Trajectories

Feasibility has financial and temporal dimensions beyond the ability of a control system to execute a production schedule. A plant may possess the engineering knowledge required to remanufacture worn equipment but lack the working capital for inventory, the credit required for tooling, or the authorized budget to maintain a line during a prolonged transition. Conversely, a profitable expansion project may be financed rapidly even when an alternative output mix would have lower physical throughput. Neither outcome follows merely from investor ownership: public budgets, cooperative reserve rules, procurement schedules, and commons-maintenance contributions can also enable or block investment. The relevant variables are access to capital, exposure to loss, liquidity over the project horizon, allocation of residual claims, and the procedures through which competing investments are evaluated.

A research design should distinguish three horizons. The operational horizon concerns the hours or days over which sensors, production orders, and controllers are adjusted. The asset-reproduction horizon concerns replacement, repair, training, software support, and the sustained availability of material inputs. The ecological and social-reproduction horizons concern changes whose burdens or recovery periods may exceed the contractual life of a facility or the tenure of its managers. These horizons are observational and institutional constructions with specified temporal boundaries; they are not treated as an ontologically fundamental partition of time. The mismatch among them matters when the benefits of rapid expansion are recognized within one accounting period while extraction, ecological exposure, or maintenance obligations extend beyond it.

The financing comparison should examine the *actual decision rule* governing each investment. A private project may be constrained by a hurdle rate, covenants, and liquidity. A public installation may use a multiyear appropriation, mandated service outcomes, and oversight. A cooperative may require member approval, reserve retention, or borrowing on terms different from those available to a corporation. A commons-oriented initiative may combine public infrastructure, contributions in kind, association funds, and shared maintenance obligations. None of these features alone identifies which project is economically efficient or normatively justified. A credible comparison records financing costs, contractual restrictions, distribution of downside risk, and the resulting maintenance or expansion decisions, with ownership held analytically separate from the measured contractual mechanism.

The capital problem also concerns reproduction of upstream capacities. An autonomous laboratory may be able to generate candidate materials experiments continuously, while the organizations that make calibrated instruments, replacement components, reagent precursors, and electrical infrastructure operate under other institutions and capital constraints. Authority over these inputs may establish effective

DRAFT

Table 43: Institutional comparison dimensions for technically comparable autonomous-production installations. Entries specify questions and plausible governance mechanisms, not observed average outcomes or universal properties of organizational types.

| Arrangement | Primary rights and financing to document | Possible choice and accountability channels | Central uncertainty for empirical research |
|---|---|---|---|
| Investor-owned | Residual claims, board and management authority, debt, contractual control of inputs | Investment return, creditor covenants, corporate governance, regulation, worker and community claims | Whether priced and non-priced ecological and labour effects alter actual decisions |
| Publicly operated | Statutory mandate, delegated management, appropriations, public procurement | Service obligations, audits, elected or administrative oversight, legal review | Whether procurement, political cycles, and performance mandates permit responsive stewardship |
| Worker or user cooperative | Membership eligibility, voting rights, retained surplus, member credit and guarantees | Member deliberation, elected management, shared risk, external regulation | Which affected persons have voice and whether finance or scale creates exclusion |
| Commons-governed installation | Access and exclusion rules, shared infrastructure, contribution and maintenance duties | Participatory rule revision, monitoring, graduated response, dispute processes | Whether resource boundaries, contributors, and affected outsiders can be governed at the required scale |

bottlenecks even when the laboratory's internal planning has become highly automated. As argued in Section 10, observing a technical dependency does not by itself establish monopolistic power, exploitative conduct, or the absence of substitutes. Those claims require records of alternatives, switching costs, access terms, and consequential decisions.

## 26.3 Comparing Institutional Arrangements on Shared Analytical Terms

An institutional comparison becomes misleading if ownership labels substitute for measurements of how an installation is governed. Hansmann's analysis of enterprise ownership provides a rationale for examining different classes of patrons and the costs of contracting and ownership, rather than treating the investor-owned corporation as the only relevant organizational form [52]. Ostrom's research on common-pool resources and polycentric governance further indicates that collective management depends on specific rules, monitoring, conflict resolution, and institutional fit; a resource does not become durably governed simply by being called a commons [146, 147, 148, 29]. Evidence surveyed by Pérotin likewise challenges simple generalizations about worker-cooperative viability while remaining dependent on sectors, countries, selection, and the outcomes examined [151]. None of these sources establishes a comparative ecological ranking for autonomous AI production.

Table reftab:institutional-comparison-types specifies observable rights and decisions rather than assigning outcomes to ownership labels. A public research facility may employ proprietary instruments and external cloud services; a cooperative may retain member voting rights while licensing a closed industrial controller; and a privately financed plant may face enforceable ecological permits and substantial worker representation. Such hybrid arrangements imply that a causal comparison should record the holders and practical exercise of particular rights, their jurisdictional basis, and the sequence of decisions that altered productive options. Matching facilities by owner category alone would confound financing, technology, regulatory context, informational access, and the institutional mechanisms under investigation. Commons governance likewise requires evidence of actual rules, monitoring, and conflict resolution, rather than a nominal designation.

Governance form should also be distinguished from productive purpose. The relevant output may serve subsistence, public health, repair of shared infrastructure, market demand, scientific inquiry, or

DRAFT

optional consumption. Equivalent throughput is not equivalent social value, and two outputs with similar economic value need not have comparable ecological functions. A comparative study should hold delivered functions or service quality constant when studying the efficiency of organizational arrangements; where the organizations deliberately select different functions, it should report that divergence as an outcome requiring normative evaluation rather than force a single unexamined denominator.

## 26.4 Authority over Automated Decisions and Infrastructural Knowledge

Advanced control systems redistribute the practical ability to make production decisions. An optimizer can recommend dispatch schedules, supplier substitutions, maintenance priorities, or production quantities. That computational recommendation acquires practical force through permissions, interfaces, staff instructions, contracts, and the possibility of intervention. The installation's institutional architecture determines who sets its objective, defines its constraints, approves updates, sees its logs, overrides a decision, and bears responsibility for an unsafe or ecologically damaging operation. Attribution of agency exclusively to "AI" obscures these arrangements, particularly when the same recommendation can be accepted, rejected, or modified by different authorized actors.

Formal and effective authority can diverge. A worker may possess a nominal right to halt a process but be unable to interpret an opaque warning or risk sanctions for stopping the line. A regulator may have a legal inspection power but lack access to relevant proprietary logs. Conversely, a specialist without the final formal vote may effectively determine the feasible options by controlling calibration, system credentials, or the description of available alternatives. Aghion and Tirole's account of formal and real authority provides an established basis for asking whether information asymmetry modifies the actual allocation of control [4]. A GR trace could record these changes if it includes both permissions and observed interventions; a list of nominal users or a software-architecture diagram alone would be insufficient.

This question has a direct engineering counterpart. An industrial control architecture typically distinguishes enterprise planning, supervisory control, equipment-level control, and safety functions. Formal safety requirements and occupational duties may restrict what even a highly optimized controller is permitted to execute; their content and enforceability depend on the technology, installation, and jurisdiction [62, 94, 176]. Researchers should document the implemented authority path through each layer, including the conditions under which a human can interrupt operation, the effect of communication failure, and the supplier's ability to revoke or change software access. The presence of a shutdown button is a design fact; whether it supplies a meaningful and safe avenue for workers or affected parties to contest a production decision is an additional institutional question.

Knowledge circulation can both extend and limit productive autonomy. Shared maintenance protocols, interoperable interface specifications, and reproducible experiments may reduce dependency on one information holder. Restrictions intended to protect security, personal data, trade secrets, or legitimate intellectual-property interests may limit disclosure. The research problem is not answered by declaring all operational knowledge open or all access restrictions illegitimate. Rather, it is necessary to determine which records enable independent evaluation, which must remain confidential, who can authorize controlled disclosure, and how errors can be corrected without creating additional cybersecurity or safety risks. W3C PROV-O supplies an ontology for representing provenance relations among entities, activities, and agents [196]. The FAIR principles separately address the findability, accessibility, interoperability, and reusability of data and metadata [191]. Neither guarantees that an asserted event occurred or that a record is scientifically adequate; a new relational layer must demonstrate additional value against these distinct baselines.

## 26.5 A Conditional Model of Institutional Implementability

The viability analysis in Section 24 considered conditions under which production and ecological constraints can be respected over a declared horizon. Institutional implementability is related to that analysis

DRAFT

but has a different logical status. Let $\mathsf{S}$ be a typed physical and ecological state space, $\mathsf{U}(s)$ the physically admissible control actions at state $s \in \mathsf{S}$, and $T \in \mathbb{N}$ a finite decision horizon measured in defined operating intervals. For a declared initial state, uncertainty model, and transition law, let $\Pi_{\text{phys}}(T)$ be the set of nonanticipative policies satisfying the chosen engineering constraints; let $\Pi_{\text{ecol}}(T) \subseteq \Pi_{\text{phys}}(T)$ be those additionally satisfying the model's stated ecological conditions. Such sets are model-relative and may change with the information structure, boundary, acceptable uncertainty, and ecological criteria.

An institutional regime $g$ is represented for this purpose by a documented collection of control and access rules, financing arrangements, resource commitments, and decision procedures over the same horizon. Let $\mathsf{Z}_g$ denote the set of admissible institutional implementation records, including expenditure approvals, procurement commitments, permissions, and required oversight decisions. Define $\mathcal{I}_g \subseteq \mathsf{Z}_g \times \Pi_{\text{phys}}(T)$ as a relation whose elements $(z, \pi)$ mean that institutional record $z$ makes policy $\pi$ implementable under the model's specified conventions. The *institutionally implementable* set is then

$$\Pi_g(T) = \left\{\pi \in \Pi_{\text{phys}}(T) : \exists z \in \mathsf{Z}_g \text{ with } (z, \pi) \in \mathcal{I}_g\right\}. \tag{71}$$

Here "implementable" is a claim about the stated institutional model, not a prediction that an authorized policy will be chosen or executed successfully. The designation does not imply that affected parties accept the institution's decisions, that every permission is lawful under all applicable rules, or that effective control and formal entitlement coincide. The policies that also meet the ecological model are $\Pi_g^{\text{ecol}}(T) = \Pi_g(T) \cap \Pi_{\text{ecol}}(T)$. Ecological feasibility and institutional permission are therefore distinct filters. The former expresses chosen physical and ecological constraints; the latter represents the modeled conditions of authorization and resource access. Where an ecological requirement is also a legal condition, the same underlying rule may appear in both descriptions for different analytical reasons; it should not be counted as two independent constraints or two independent sources of warrant. Neither set resolves the legitimacy of those filters.

**Proposition 17** (Conditional containment and institutional feasibility)**.** *For fixed $T$, physical model, initial state, and definition of institutional implementation, suppose regimes $g_1$ and $g_2$ have implementation relations satisfying*

$$\text{proj}_\Pi \mathcal{I}_{g_1} \subseteq \text{proj}_\Pi \mathcal{I}_{g_2}, \tag{72}$$

*where* $\text{proj}_\Pi$ *denotes the projection onto policies. Then* $\Pi_{g_1}(T) \subseteq \Pi_{g_2}(T)$ *and* $\Pi_{g_1}^{\text{ecol}}(T) \subseteq \Pi_{g_2}^{\text{ecol}}(T)$.

*Proof.* By Equation (71), $\Pi_g(T)$ is the policy projection of $\mathcal{I}_g$ restricted to $\Pi_{\text{phys}}(T)$. The assumed projection inclusion gives the first set inclusion. Intersecting both sets with the fixed set $\Pi_{\text{ecol}}(T)$ preserves inclusion, giving the second. No ranking of ecological performance among selected policies follows from either inclusion. For example, with identical physical conditions and a burden indicator measured in kg per fixed service period, let $g_1$ implement only policy $A$ with burden $10$ kg, while $g_2$ can implement both $A$ and policy $B$ with burden $14$ kg. The inclusion holds, yet selecting $B$ under $g_2$ increases the recorded burden. These are constructed quantities, not evidence about ownership forms. A policy-selection mechanism and verified consequences are required to derive an ecological comparison from an institutional capability relation. □

The proposition is a modest set-theoretic implication of explicitly assumed implementability relations. It is *not* evidence that private, public, cooperative, or commons governance generally supplies more freedom, and its inclusion premise cannot be inferred from a regime's title. In particular, expanding formal access may affect congestion, incentives, operating knowledge, investment, or physical conditions; the fixed-physical-model assumption then fails. Nor does a larger implementable set require decision makers to choose a lower-burden policy. For example, suppose two policies are technically and ecologically admissible under declared bounds, but the newly available one has higher energy use than the previously chosen policy. A regime that enlarges the choice set can select the higher-use policy. The physical and normative evaluation therefore depends on selection rules and actual outcomes rather than the cardinality of authorized alternatives.

DRAFT

An empirical implementation must specify how $\mathcal{I}_g$ is identified. A suitable record may require measured available credit, signed supply agreements, observable authority to alter control limits, and evidence of safe implementation capacity. A missing contract need not prove a policy impossible if an alternative supplier or procedure is available. Similarly, formal permission is not sufficient if staff cannot safely execute the policy. Such qualifications are especially important when institutional authority is distributed across a supply network rather than concentrated in a facility's legal owner.

## 26.6 Productive Power, Ecological Burdens, and the Claims of Others

A physical inventory records flows and exposure within chosen boundaries; it does not itself identify who has the authority to generate them, who receives the resulting service, or who has an enforceable claim against their adverse consequences. The distinction matters when a production choice produces benefits within a firm but requires extraction in another region, uses shared water resources, or exposes workers and residents to risks not captured by the firm's accounts. The literature on social-ecological governance emphasizes that ecological problems often cross the organizational or local boundaries within which rules were formed [29, 147]. Lifecycle assessment and multiregional input–output models can trace some displaced resource demands, but they do not automatically establish legal responsibility or procedural legitimacy [83, 190].

Control over a technical bottleneck may allow one actor to determine which downstream users gain access to production capacity or process knowledge. Yet unequal bargaining power, market concentration, legal authority, and socially unjust extraction remain distinct propositions. Evidence should identify actual restrictions, available alternatives, switching costs, the incidence of expenses and benefits, and the institutional avenues through which affected parties can object. Contract terms or accounting margins alone cannot establish the experienced welfare of workers, the preservation of ecological processes, or the legitimacy of an allocation. This distinction also prevents a misleading inference in the other direction: observing that an installation complies with local legal duties does not imply that its cumulative effects beyond the regulatory boundary are acceptable.

Public and collective institutions face their own distributional questions. Voters do not always coincide with affected communities; cooperative members may exclude temporary workers, consumers, or upstream suppliers; and a commons association may place maintenance obligations on participants without giving outside ecological interests representation. A credible comparison consequently records both the membership of decision-making bodies and the population bearing consequences. It should distinguish the capacity to be informed, the capacity to participate, the authority to alter a decision, and access to appeal or remedy. These rights may be established through different laws and institutions in different jurisdictions; the present section proposes comparison variables rather than claims about universal legal entitlements.

For ecological decisions involving nonhuman life, the representative problem requires particular care. Ecosystems cannot be assumed to be adequately represented by a firm's projected costs or by the preferences of a single class of owners. Selection of ecological indicators, reference conditions, acceptable disturbance, and recovery targets involves scientific uncertainty and normative judgments. Work on plural environmental values supplies reasons to make those interpretive choices explicit [60]. Scientific measurement can constrain false claims and expose incompatible objectives; it cannot independently authorize a particular constituency to sacrifice another's interests. The IPBES values assessment evaluates plural approaches to valuation and decision-making; it neither assigns legal standing to particular nonhuman entities nor validates any one institutional representation procedure for an autonomous factory. This difference between evidence and authority is a fundamental condition for interpreting the conditional ecological sets above.

## 26.7 A Constructed Comparison with Identical Productive Equipment

Consider four hypothetical installations, each assumed to possess an identical automated fabrication cell, the same verified maintenance capability, and access to the same technically qualified parts and elec-

DRAFT

Table 44: Hypothetical audit of barriers to executing the same technically available policies. All entries are constructed assumptions; no empirical superiority is implied.

| Installation | Assumed constraint for $\pi^R$ | Record required to verify an implementability claim | Possible change to the finding |
|---|---|---|---|
| Investor-owned | Lender approval for tooling expenditure | Debt terms; consent record; tooling invoice | Different credit terms remove the barrier |
| Publicly operated | Mandate approval for a revised service mix | Procurement plan; delegation and service mandate | Revised mandate permits execution |
| Worker cooperative | Member approval and qualified training | Voting rule; minutes; competence records | Prior delegation allows faster action |
| Commons-governed | Licence access to external calibration software | Access rule; licence; alternative model validation | Open alternative eliminates dependency |

tricity infrastructure. Each cell can supply a specified functional output under a standard schedule or shift some capacity toward component remanufacturing. All four have an identical physical engineering model and ecological measurement protocol; their financial commitments, access rights, governance procedures, and effective control over the controller differ. The example is a mechanism-isolation exercise. It represents no observed installation and contains no estimated treatment effect.

Suppose a standard operating policy $\pi^S$ yields 100 functionally comparable service units over the declared accounting interval and a repair-oriented policy $\pi^R$ yields 90 new service units plus 15 independently qualified repaired components. These outputs are deliberately *not* added into a single service total: repaired components and new units have different service definitions, and their lifecycle effects require separate functional equivalence tests. Assume both policies are physically available. A third, higher-throughput policy $\pi^E$ requires an additional supply commitment and has greater counted electricity demand under the narrow operating boundary. The engineering facts are fixed for this illustration, while any ecological ranking remains contingent on actual embodied flows, replacement avoidance, output demand, and site-specific exposure. The choice can therefore be studied without assuming that the repair-oriented policy is globally preferable.

In an investor-owned installation, suppose management can immediately authorize $\pi^S$ but financing for remanufacturing tooling depends on a lender's consent. In a publicly operated installation, suppose remanufacturing is funded but changing the delivered-service mix requires approval under the service mandate. In a worker cooperative, suppose both policies are technically financed but require a membership vote and a contracted assurance of worker training. In a commons-governed facility, suppose remanufacturing is permitted by shared operating rules, while use of an externally maintained calibration model requires a separate licence. These are invented institutional conditions, not essential characteristics of the forms. Each regime can make a different subset of the same physical policies immediately implementable; different financing or contractual terms could reverse the result.

The comparison clarifies three separate inferences. First, technically identical cells may encounter different conditions of institutional implementability. Second, one can test the causal significance of a specific barrier, such as a loan covenant, only if the barrier and alternative policies are measured with sufficient independence from productive performance. Third, a favourable outcome in one example would neither establish the general superiority of its institutional form nor show that its procedure is legitimate for people outside the decision-making body. A strong empirical study would also consider differences in managerial capacity, maintenance culture, market exposure, public subsidies, and preexisting access to expertise. These potential confounders cannot simply be standardized away by matching robot models.

## 26.8 Identification, Comparative Baselines, and Potential Negative Results

The empirical question is whether a specified institutional control mechanism changes the feasible and selected production decisions, and with what measured consequences, after accounting for material and ecological conditions. An investigator should avoid treating an ownership dummy as a complete inter-

DRAFT

Table 45: Prospective empirical tests separating institutional mechanisms from organizational labels. "Negative result" denotes evidence that would weaken a proposed mechanism, not a statistical conclusion already established.

| Proposed mechanism | Data required | Competent baseline or design | Informative contrary result |
|---|---|---|---|
| Finance restricts repair | Credit terms, tooling cost, repair demand, qualifying failures | Within-facility change; matched credit access | Repair unchanged after constraint lifted |
| Access rights restrict substitution | Licence scope, alternative qualifications, switching cost, procurement histories | Conventional supplier-risk analysis | No material dependence on restricted provider |
| Decision rights affect scale | Authorization logs, demand, prices, service, energy/material flows | Matched policies; explicit demand control | Decisions insensitive to rule changes |
| Disclosure changes accountability | Audit access, correction time, independent error detection, remedies | Existing logs, LCA and provenance systems | Added trace offers no improvement |
| Ecological safeguards alter burden | Compliance records, absolute flows, upstream transfer, ecological receptors | Lifecycle and cumulative-impact assessment | Local gain offsets elsewhere or no attributable improvement |
| Distribution of benefits changes | Worker outcomes, access to service, costs and risk incidence | Labour and distributional survey with controls | Differences explained by wages, sector, or selection |

vention. More informative exposures include a documented change in authorization rules, the entry or expiry of a software licence, a shift in the financing terms for repair, a revision to worker participation, or introduction of independently reviewable ecological reporting. Where such changes occur within a continuing installation, an interrupted-series or event-study design may be informative, but only if there is a defensible comparison and competing changes are considered. Cross-sectional matching can describe differences; it seldom removes unobserved heterogeneity in management, initial capital stocks, demand, local regulation, or selection into ownership forms.

The unit of analysis should follow the mechanism. A software-licence restriction is observed at the controller or contracting relationship; replacement investment may require facility-level records; regional emissions and ecological effects require geographically defined material-flow and receptor data. For a credible institutional comparison, researchers should retain the service definition, observation horizon, technical configuration, and measurement procedures wherever scientifically defensible. They should report deviations rather than claim exact matching where supply networks, labour arrangements, or demand conditions are inherently different. Spillovers are especially important because one facility's access to a scarce input may alter another's availability, and a single supplier may govern several nominally independent plants.

An evidence protocol should preserve dated decisions, actors' documented authority, the versions of models and technical constraints available at each decision, refused alternatives, and subsequent outcomes. Some materials may be confidential or personally sensitive; their treatment requires controlled access, aggregation, or independent review rather than unqualified public disclosure. Data provenance should identify what was measured directly, reconstructed from records, estimated through an economic model, or asserted in interviews. Public reporting should differentiate verified violations, contested claims, uncertainty in burden attribution, and unresolved normative disagreement. An institution-specific case study may support a mechanism without identifying a transferable average effect across organizational types.

The proposed GR representation faces a concrete burden of proof. If existing enterprise-resource planning logs, production-network data, formal governance documents, environmental accounting, and established provenance models reconstruct the relevant decision pathways with equal or better fidelity and cost, the additional GR layer has no demonstrated methodological advantage. If decision rights and technical dependencies vary too rapidly or are withheld, a detailed relational graph may convey false

DRAFT

precision. Equally, where ecological indicators are too sparse to discriminate policies, a richer account of authorization cannot compensate for missing ecological evidence. These are possible negative findings to record, not implementation failures to hide.

### 26.9 Implications for Comparative Evaluation of Production Policies

The section has established a conditional conceptual distinction among physical production capability, institutionally implementable action, observed policy selection, and legitimate authority to choose. Property, financing, information, procurement, workplace relations, public obligations, and dispute procedures can constrain or expand which production trajectories are available in practice. The formal result is an elementary inclusion statement under explicitly assumed relations; the constructed installation comparison illustrates mechanisms but supplies no empirical estimate. Existing work in property-rights theory, organizational authority, enterprise ownership, and commons governance already addresses important parts of this argument, and the suggested GR extension remains subject to comparative validation.

For historical-materialist inquiry, the central research issue is how transformations of machinery and technical knowledge interact with the reproduction of the institutions that command them. An autonomous controller may reduce the need for continuous direct intervention while increasing reliance on licensors, lenders, specialist maintenance providers, data infrastructures, or concentrated equipment suppliers. It may also support distributed repair, shared learning, and more contestable decisions under suitable institutional conditions. These outcomes cannot be ranked in advance by ownership designation. Their evaluation requires material accounting, measurements of labour and ecological effects, and evidence of who can shape and challenge the decisions through which productive possibilities are realized.

Section 27 therefore proceeds to counterfactual evaluation of production policies. Its comparisons should specify both engineering and institutional interventions, distinguish policy assignment from implemented behaviour, and examine throughput-, energy-, and ecology-oriented choices with credible conventional baselines. Section 28 will consider the observable data needed to measure these mechanisms in concrete cases, while Section 29 will evaluate the resulting limits of causal and normative inference. No empirical conclusion about the ecological superiority of any ownership form is claimed here.

## 27 Counterfactual Evaluation of Production Policies

Section 26 distinguished physically feasible production from decisions that are institutionally implementable, selected, and justified. This section establishes how competing production policies could be compared without treating a computational optimum, a favorable lifecycle inventory, or an ownership category as evidence of a causal ecological improvement. The immediate research task is to specify a reproducible counterfactual design: what service is supplied, which policy is changed, what would otherwise have happened, whose burdens are measured, and what findings would contradict the proposed account. The comparison places throughput-oriented, direct-energy-oriented, and ecologically constrained controllers alongside a competent conventional baseline. It addresses both controlled operating experiments and more difficult observational comparisons, and it makes room for interactions among facilities, suppliers, electricity systems, and affected ecological receptors. No original trial, simulated result, or estimated causal effect is claimed here. The proposed methodology builds on established experimental design, causal inference, manufacturing measurement, and lifecycle assessment; the generative-relational (GR) representation is evaluated only for the additional information it might provide.

### 27.1 Defining the Service that Policies Must Supply

A fair experiment begins before assigning a controller. The investigators define a *productive service*, not simply a raw item count: an eligible product or experimental output with declared composition,

DRAFT

tolerance, reliability, safety characteristics, delivery window, and subsequent service life where appropriate. A polymer component that passes the same dimensional inspection but fails much earlier in use cannot count as an equivalent output without accounting for the required replacement. Likewise, a laboratory's number of reactions is not a suitable common denominator if the competing policies have different success criteria or scientific information content. ISO 14040 and ISO 14044 provide a pre-existing framework for goal and scope definition, functional units, reference flows, inventory analysis, interpretation, and limitations in lifecycle comparisons [84, 83]. Manufacturing KPI conventions are available in ISO 22400-2, although no single factory KPI represents ecological performance [86].

Let $i \in \mathcal{I}$ index eligible production cells or facilities, $t = 0, \ldots, H$ index prespecified operating intervals, and $d_i = (q_i, \kappa_i, \tau_i)$ denote the assigned service requirement. Here $q_i$ is a quantity in the study's functional units (FU), $\kappa_i$ is a vector of independently measured quality and safety criteria, and $\tau_i$ gives required delivery or performance windows. The horizon $H$ is measured in explicitly stated calendar or operating intervals, and eligibility is established from *pre-assignment* information. Let $\mathcal{D}_i$ represent allowed variations in demand and input specification, fixed or randomized before policy assignment. A conforming-output function $Q_i(\pi, w)$ counts delivered FU under policy $\pi$ and operating conditions $w$ only when all predetermined quality gates are met. A proposed matched-service requirement is

$$\mathcal{P}_i^{\text{svc}} = \left\{ \pi \in \mathcal{P}_i : Q_i(\pi, w) \geq q_i \text{ and } \kappa_i(\pi, w) \in \mathcal{K}_i^{\text{quality}} \quad \forall w \in \mathcal{W}_i \right\}, \tag{73}$$

Here $\mathcal{W}_i$ is a predeclared set of operating conditions, and the displayed universal quantifier expresses a robust criterion over that set. A finite benchmark set, a validated uncertainty envelope, and a probabilistic service guarantee encode different substantive assumptions and must not be interchanged; a chance-constraint design would require a separate formula and probability measure. In particular, quality measured *after* treatment is an outcome, and restricting analysis to the jobs that subsequently passed inspection can induce selection bias. For the primary comparison, failures to meet the service are recorded and reported as failures rather than silently discarded or normalized away.

An observation needs at least two service perspectives. The *fixed-service estimand* compares operating strategies at the predeclared $q_i$ and $\kappa_i$. It can identify an engineering opportunity to reduce inputs for a comparable output if identification conditions are met. The *system-response estimand* permits output demand, price, capacity expansion, adoption, and repair or replacement patterns to respond to the policy. It asks whether aggregate resource requirements or exposures actually change. Section 25 established why the two questions can have opposite answers. Neither is a substitute for the other. If environmental constraints make $q_i$ infeasible at a facility, the study should report an empty feasible set and examine alternative service levels or locations in a separately labeled decision analysis rather than redefine the original treatment's success.

## 27.2 Specifying the Interventions and Competent Comparators

The treatment is an executable *policy protocol*, including its objective, constraints, information access, allowed interventions, hardware, human roles, and procedures for safe exceptions. A label such as "AI optimization" is insufficient: it may refer to schedule adjustment, equipment replacement, predictive maintenance, a new recipe, or demand management, each with different causal consequences. Denote the four candidate strategies by $a \in \mathcal{A} = \{0, \mathrm{T}, \mathrm{E}, \mathrm{C}\}$. Strategy $0$ is a competent incumbent industrial operating method, including its existing automation, safeguards, maintenance expertise, and economically reasonable tuning. Strategy T prioritizes productive throughput or schedule performance while respecting identical mandatory safety and quality rules. Strategy E prioritizes *direct metered energy* for matched service, subject to the same mandatory rules. Strategy C seeks policies satisfying separately justified ecological constraints, including specified upstream and downstream burdens where the necessary data are available. The ecological protocol's optimization within its feasible region must be documented rather than inferred from a favorable outcome after the fact.

For formal clarity, let $L_i^{\text{svc}}(\pi)$ denote service shortfall, $R_i(\pi)$ represent a turnaround or throughput objective in declared units, $E_i^{\text{meter}}(\pi)$ be electricity plus separately reported on-site fuels in physical

DRAFT

Table 46: Proposed policy contrasts under an identical service specification. Entries are prospective design requirements, not observations or average treatment effects. All arms retain mandatory legal, safety, and quality requirements.

| Arm | Governing operating rule | What must be held comparable or recorded | Principal comparison question |
|---|---|---|---|
| 0 | Competent incumbent workflow | Same service, load, equipment state, staffing qualification, and safety standards; incumbent tuning and maintenance documented | Does any intervention outperform an existing defensible practice? |
| T | Prioritize schedule or throughput performance | Service floor, downtime, run length, scrap, overtime, equipment wear, and exceptional interventions | What burdens accompany additional production flexibility? |
| E | Minimize direct measured energy per matched service | Meter boundary, electrical load profile, fuel, idle loads, quality, upstream energy and changed cycle times | Does reduced direct energy persist across the lifecycle and response boundary? |
| C | Respect predeclared ecological limits before optimizing an operating criterion | Receptor-specific constraints, material and lifecycle inventories, data gaps, quality, and feasibility failures | Can specified service be delivered while respecting defensible ecological conditions? |

energy units after a documented conversion convention, and $\mathbf{B}_i(\pi) = (B_{i1}, \ldots, B_{iK})$ be a vector of lifecycle inventories or impact-category results with each component retaining its own unit. The three exploratory policy designs can be written schematically as

$$\begin{aligned} \pi_i^{\mathrm{T}} &\in \arg\min_{\pi \in \mathcal{P}_i^{\mathrm{svc}}} R_i(\pi), \\ \pi_i^{\mathrm{E}} &\in \arg\min_{\pi \in \mathcal{P}_i^{\mathrm{svc}}} E_i^{\mathrm{meter}}(\pi), \\ \pi_i^{\mathrm{C}} &\in \arg\min_{\pi \in \mathcal{P}_i^{\mathrm{svc}} \cap \mathcal{P}_i^{\mathrm{ecol}}} C_i^{\mathrm{decl}}(\pi), \qquad \mathcal{P}_i^{\mathrm{ecol}} = \{\pi : \mathbf{B}_i(\pi) \preceq \mathbf{b}_i\}, \end{aligned} \tag{74}$$

The displayed arg min expressions specify selected policies only when their respective feasible sets are nonempty and their objective infima are attained (for example, under appropriate compactness and lower-semicontinuity conditions). If not, the study must report infeasibility or use a prespecified approximate-optimization criterion with its tolerance, rather than silently asserting the existence of an optimal controller. Here $C_i^{\mathrm{decl}}$ is a declared monetary, reliability, or other operational criterion, and $\mathbf{b}_i$ contains category-specific bounds with matching units and explicitly identified scientific and institutional authority. The symbol $\preceq$ means componentwise comparison only where the indicators are defined on compatible scales. The rightmost condition describes a conditional *model*, not proof that the relevant ecological thresholds have been identified or that the reported categories exhaust ecological harm. Other policy objectives, including lexicographic thresholds or Pareto-efficient designs, are possible. Comparing alternatives with different mandatory constraints requires disclosing this difference as part of the treatment rather than presenting it as a purely algorithmic performance gain.

The baselines must remain credible. A poorly tuned manual line is not a sufficient control for a commercially competent automated cell; equally, an expensive ecological controller should not be credited with the effects of new meters or repaired machinery that could also have been fitted to the incumbent system. A factorial design can separate controller selection from capital upgrades, supply contracts, or additional monitoring. Where a strategy relies on additional sensing, cloud inference, model training, or human supervision, its energy, material, labor, downtime, and expense enter the treatment record. Safety-related stop actions and operator overrides are logged and evaluated for adherence and risk; they cannot be hidden as exceptions to make a policy appear independently autonomous.

DRAFT

## 27.3 Causal Questions and the Assignment Mechanism

A counterfactual is an outcome under an alternative course of action in a sufficiently specified setting. For fixed-service experiments, define $Y_{ik}(a; H, \mathcal{B})$ as the potential cumulative outcome for eligible unit $i$, outcome category $k$, policy arm $a$, horizon $H$, and physical/accounting boundary $\mathcal{B}$. Each $Y_{ik}$ has its declared unit: kilowatt-hours, kilograms of particular material, kilograms of $CO_2$ equivalent under a specified characterization method, cubic metres of consumed water, worker-hours, or a receptor-specific ecological quantity. Potential outcomes are conceptual targets; an investigator observes only the realized arm under a particular experimental arrangement. A predeclared average contrast, among units eligible before assignment, is

$$\Delta_k(a, 0 \mid \mathcal{E}, H, \mathcal{B}) = \mathbb{E}[Y_{ik}(a; H, \mathcal{B}) - Y_{ik}(0; H, \mathcal{B}) \mid i \in \mathcal{E}], \tag{75}$$

where $\mathcal{E}$ is the population of eligible units and the expectation refers to its specified sampling or assignment distribution. An average improvement in one $k$ gives no automatic ordering for other components, and an average can conceal important exposed subpopulations. A strong experimental report provides both distributions and explicitly warranted thresholds rather than one unexplained weighted "sustainability score." The general discipline of specifying a target trial's eligibility, treatment, time zero, outcomes, assignment, and analysis plan is well developed in causal inference [54]. That methodological framework does not itself establish that exchangeability, positivity, consistency, or adequate measurement hold in a particular production dataset.

Where distinct cells can genuinely be operated independently, random assignment of eligible cells or predeclared operating blocks to policy arms is a powerful first choice. Blocking by machine type, shift, product family, input grade, or season may improve precision and protect balance, provided the blocks are determined before assignment. Repeated crossovers can be economical but are suitable only when carryover is bounded: machine wear, contaminant accumulation, recalibration, personnel learning, and preventive maintenance can make yesterday's treatment relevant to today's outcome. Required washout or restoration protocols should be justified with engineering evidence. If washout is impossible, longer parallel arms or cluster-level randomized rollouts may be preferable to within-cell crossover.

For observational policy changes, the study should explicitly *emulate* the desired experiment. The analogy to target-trial emulation transfers a methodological protocol, not any biomedical effect estimate to industrial systems. Eligibility, policy assignment and follow-up must be aligned at a defensible time zero; a site that enters the record only after successfully completing an AI-assisted production campaign should not retrospectively count as if it had been at risk of recorded failure throughout that campaign. Policy adoption may depend on prior energy prices, profitability, maintenance deficits, supplier availability, labor relations, ecological regulation, or a manager's private forecasts. Matching on observed covariates, propensity-score methods, outcome regression, or doubly robust estimators can address measured confounding under their assumptions, yet do not recover unmeasured assignment determinants [161, 54]. If a policy is introduced at different dates across plants, difference-in-differences estimands should accommodate heterogeneous treatment effects and adoption timing rather than assume that a conventional two-way fixed-effects coefficient is the desired effect [17]. The particular group-time contrasts being aggregated must be declared. Parallel-trend requirements, anticipation, overlapping interventions, and absence of differential concurrent shocks need substantive scrutiny. Pre-policy outcomes, procurement records, regulatory changes, and negative-control outcomes are valuable but do not guarantee identification.

A special danger is *conditioning on a treatment-induced variable*. Suppose C lengthens cycle time and changes which orders are accepted, or T increases orders through a lower price. Comparing only completed orders of identical realized size after adoption can select unlike populations. The primary fixed-service trial therefore holds the *assigned* service schedule comparable and records delivery failures. A secondary demand-response study explicitly includes the selection, pricing, and adoption mechanism. These are different target interventions and should be reported as such. A finite-horizon effect also must not be generalized to a permanent policy regime without a defended model of capital replacement, learning, and ecological recovery.

DRAFT

Table 47: Measurement boundary and provenance for proposed fixed-service and system-response comparisons. No numeric effects are imputed; selection of particular indicators and receptors is site dependent.

| Domain | Primary measurement with units | Boundary and timing requirements | Essential limitations |
|---|---|---|---|
| Productive service | Qualified FU delivered, rejects (FU), lead time (h), uptime (h) | Assigned order and quality criteria; fixed operating horizon; customer service life where relevant | Realized completion is treatment-sensitive; quality thresholds need independent validation |
| Energy and equipment | Electricity (kWh), fuel (MJ or unit-specific), idle load (kW), equipment mass (kg), maintenance (h) | Meter timestamps, upstream equipment manufacturing, capacity and replacement history | Direct energy differs from characterized lifecycle impacts |
| Materials and water | Material inputs/outputs (kg by substance), water withdrawal/consumption ($m^3$) | Supplier origin, catchment, seasonal conditions, recovery routes, allocation rules | Missing tiers, water stress and material substitution affect comparability |
| Ecological exposure | Pollutant loads (kg by substance), receptor exposure (specified unit), habitat or biotic indicators | Site and off-site receptors, lag structure, cumulative exposures, baseline and recovery observations | No universal conversion of heterogeneous ecological processes |
| Social and institutional reproduction | Paid/unpaid labor (h), operator interventions (count), control permissions (documented), worker exposure (specified unit) | Employment and supplier tiers, authority to override, procurement, grievance records | Hours, rights, welfare, and legitimacy remain analytically distinct |

## 27.4 Accounting for the Consequences beyond the Production Cell

A policy can improve direct production metrics while moving burdens to suppliers or later lifecycle stages. ISO-style LCA, material-flow accounting, and input–output approaches provide the competent comparison baseline for locating and characterizing these burdens [84, 83, 195, 190]. For each arm, investigators should retain separate physical records for purchased electricity by time and source, fuels, heat, water withdrawal and consumption, process chemicals, conforming and rejected materials, transport, replacement equipment, computation, packaging, service life, and disposal or recovery pathways. Whether a flow is assigned to a product, facility, or supplier must be documented consistently, including co-products, shared equipment, end-of-life allocation, and uncertainty in background datasets. An attributional inventory describes the modeled burden assigned to the compared services; its difference does not necessarily estimate the *consequences* of scaling a policy when suppliers, electricity dispatch, prices, and demand change. Consequential LCA and economic response modelling can instead target marginal changes but introduce further scenario dependence and uncertainty [155].

Let $\mathbf{F}_i(a;\mathcal{B})$ be an inventory vector of typed elementary flows attributed to arm $a$ within boundary $\mathcal{B}$. For each impact category $k$, let $\chi_k$ denote a declared and versioned characterization operator with compatible units and spatial/temporal applicability. A conventional characterized indicator has the form

$$B_{ik}^{\mathrm{attr}}(a;\mathcal{B}) = \chi_k\big(\mathbf{F}_i(a;\mathcal{B})\big), \qquad \Delta B_{ik}^{\mathrm{attr}} = B_{ik}^{\mathrm{attr}}(a;\mathcal{B}) - B_{ik}^{\mathrm{attr}}(0;\mathcal{B}). \tag{76}$$

The notation does not make $\chi_k$ valid for every flow, location, or receptor. Characterization results may depend on global warming potential horizon, water scarcity characterization, toxicity fate/exposure models, or other explicitly selected conventions. For temporally resolved electricity, both kWh and the location- and interval-specific generation consequences should be examined where data permit; direct energy minimization does not entail minimum greenhouse-gas emissions. Ecological states, species responses, and exposure histories may require distinct field measures rather than numerical conversion into one lifecycle score. The difference $\Delta B_{ik}^{\mathrm{attr}}$ is an accounting contrast under an agreed rule; equivalence to $\Delta_k$ in Equation (75) requires further causal assumptions.

DRAFT

Absolute impacts should be reported with per-FU intensities. A ten-percent modeled decline per FU, for example, is not an observed ecological benefit if the total FU changes and upstream burdens are omitted. Spatially differentiated water consumption and habitat disturbance may be decisive even where aggregate emissions decline. The relevant reference scenario includes possible changes to equipment life, scrap recovery, service replacement, and capacity growth. If the ecological policy requires a rare input whose extraction is poorly documented, the correct finding may be an *undetermined comparison* rather than a favorable result. A further challenge is the treatment's effect on research knowledge itself: an autonomous laboratory that generates many low-value measurements cannot be treated as functionally equivalent to one that produces fewer, independently validated findings merely through a unit count. A scientific service requires predeclared epistemic quality criteria appropriate to that domain.

## 27.5 Interference, Market Responses, and Historical Scale

The independent-unit contrast in Equation (75) can fail to represent the desired intervention when one facility's policy changes another's outcomes. Electricity dispatch, competition for reclaimed material, upstream suppliers' production schedules, congested transport, regional water availability, and market prices generate plausible forms of *interference*. Lowering demand at one plant could alter grid generation only slightly, while widespread adoption could change investment in generation and transmission. A high-throughput controller could displace output from another factory, change total demand, or increase extraction in a shared catchment. Such effects cannot be recovered from a trial assuming isolated treatment units unless their cross-unit pathways are small under a defended exposure model.

Let $\mathbf{a} = (a_1, \ldots, a_n)$ be the assignment vector across facilities and $z$ a declared external policy and market regime. Define the network potential outcome $Y_{ik}(\mathbf{a}; z, H)$ and regional aggregate $Y_{\mathcal{R}k}(\mathbf{a}; z, H)$ on an explicitly bounded region $\mathcal{R}$. One possible programme effect of adopting arm $a$ at the eligible scale, compared with a baseline allocation $\mathbf{0}$, is

$$\Delta_{\mathcal{R}k}^{\text{programme}}(a, 0; z, H) = \mathbb{E}\Big[Y_{\mathcal{R}k}(\mathbf{a}^{\text{adopt}}; z, H) - Y_{\mathcal{R}k}(\mathbf{0}; z, H)\Big], \tag{77}$$

where $\mathbf{a}^{\text{adopt}}$ describes a prespecified adoption coverage, eligible facilities, and timing. This target differs from the effect of assigning one cell to $a$. The effect of national or international adoption depends on trade, investment, substitution, electricity systems, rebound, and cumulative ecological conditions beyond a single study facility. Network or cluster-level randomization, spatial exposure mapping, staged adoption studies, and transparent structural scenario models may illuminate some of these channels, each with significant assumptions and possible non-identification. The adopted vector is not ontologically privileged: changing the selected nodes, edge types, and interaction horizon may change the result.

Scale also creates conflicts among social groups and institutions. A policy selected by a firm may lower on-site costs and raise supply prices for another user, or concentrate burdens where ecological data and political representation are weak. Section 26 supplied the institutional variables that allow such mechanisms to be traced; nevertheless, observational variation in governance cannot be interpreted as randomized ownership assignment. A defensible design tests a specified governance mechanism—for example, mandated disclosure of input sources, a change in controller override authority, or access to a common maintenance repository—within a suitably comparable setting. Public, investor-owned, cooperative, and commons-oriented organizations may differ jointly in location, investment, market position, regulation, and monitoring quality. Directly contrasting their mean emissions without adjustment and case knowledge cannot identify an institutional treatment effect.

## 27.6 An Explicit Protocol for Implementable Comparisons

The practical protocol must separate experimental design choices from results that have not yet been obtained. Table 48 describes a prospective, adaptable factory-cell trial; its intervals and replication are to be established by engineering feasibility, expected variance, resource constraints, and an explicit precision or power analysis. This is a design illustration, not a report of enrolled sites. A broader interfacility

DRAFT

Table 48: Prospective counterfactual study specification for comparable automated production cells. "Required" identifies a protocol commitment, not already available data or a completed registration.

| Design component | Proposed requirement and interpretation |
|---|---|
| Eligibility and time zero | Declare facility capability, calibrated meters, product specification, safety rules, pre-policy condition, and assignment moment; exclude only by prespecified baseline criteria. |
| Treatment and comparison | Randomize or justify assignment of 0, T, E, and C; archive controller version, inputs, compute and equipment additions, human supervision, exception rules, and institutional permissions. |
| Experimental unit and replication | Use independent cells or interference-aware clusters; block by pre-treatment characteristics. Determine replicate counts from feasibility and uncertainty, not an invented sample-size claim. |
| Primary outcome and estimand | Prespecify a category-specific absolute inventory or direct physical burden over the same assigned service and horizon; report service failures separately, together with all co-primary categories. |
| Secondary outcomes | Per-FU intensities; conforming output; electricity time profile; equipment wear, maintenance, material substitution, labor and safety; independently measured ecological indicators when feasible. |
| Identification and statistical plan | Publish assignment code, baseline balance checks, missing-data plan, crossovers and carryover treatment, confidence/uncertainty intervals, and sensitivity analyses; report deviations. |
| Lifecycle extension | Describe inventory boundary, process and supplier data, background database version, allocation method, characterization version, unresolved flows, and alternative demand scenarios. |
| Governance and ethics | State legal permissions, safeguards, operator authority and pay arrangements, supplier confidentiality limits, consultation with affected groups, and procedure for unsafe or ecologically harmful outcomes. |

analysis of expansion and upstream ecological consequences is a distinct study with a separate preregistration and identification argument. The distinction avoids overstating what a carefully controlled factory experiment can show about a regional economy.

A primary analysis can report assignment effects by randomized arm, retaining units that miss the target because such failures bear on feasible implementation. A secondary per-protocol analysis may examine realized adherence, but adherence itself can depend on initial assignment, input prices, breakdowns, and prior outcomes. Conditioning naively on compliant runs can therefore distort the causal comparison. Where policy regimes evolve during the horizon, longitudinal methods with explicit time-varying covariates may be required [54]. More generally, investigators should report estimation failures when feasible-policy overlap disappears, instruments fail calibration, a control law is unsafe, or substantive outcomes are systematically unobserved. These are informative findings, not merely reasons to remove adverse observations.

The protocol also requires honest measurement uncertainty. Calibrations, sampling intervals, sensor drift, missing suppliers, chemical composition, background inventory models, and ecological response functions introduce different kinds of uncertainty. Laboratory instrument uncertainty can be described using established metrological conventions, including the JCGM guide; uncertainty from an absent upstream supply network or contested ecological threshold is not exhausted by a sensor error bar [107]. Where appropriate, investigators should propagate parameter uncertainty through inventories and outcome models, distinguish variance due to assignment from measurement and model uncertainty, and report structural scenario sensitivity rather than collapse it into an unjustified single confidence interval. Data access and confidentiality restrictions should be recorded in a provenance ledger; released aggregate estimates alone may be insufficient for independent reanalysis.

## 27.7 Evaluating the Additional Information Supplied by Relational Representation

Industrial ecology, process engineering, optimization, established production-network analysis, and causal inference already provide substantial means of assessing policy outcomes. The proposed GR description should consequently be tested against strong alternatives, not against a deliberately impoverished inventory spreadsheet. Its candidate role is to preserve provenance-sensitive trajectories of decisions, dependencies, repair, authority, and exposures across a production network when important

DRAFT

relations change. A possible contribution would be to identify that a direct-energy saving is purchased through an unobserved upstream input, or that an ecological safeguard is formally configured but cannot be enacted because a supplier or software vendor controls the relevant interface. Such explanations are hypotheses requiring comparison with conventional multi-tier bill-of-materials analyses, supply-chain maps, lifecycle inventories, and temporal process models.

One evaluation can compare otherwise identical inference tasks using (i) an established LCA/MFA and ordinary causal-analysis workflow, (ii) a competent temporally annotated production-network model, and (iii) that same workflow augmented with typed GR provenance and institutional-event records. Blind independent analysts could be asked to identify missing dependencies, recover the provenance of consequential decisions, predict which recorded shocks cause service failure, or find contradictions between an ecological commitment and an executable policy. The preregistered outcomes would include precision and recall against adjudicated ground-truth cases, prediction error with suitable time splits, uncertainty calibration, false-positive burden, analyst effort, data-collection cost, compute, and reproducibility. Training and test cases must be separated by production history or facility when shared templates would otherwise leak answers. The GR mechanism should be considered unsupported where ablation shows no reproducible improvement, where advantages vanish after giving the conventional network model the same information, or where the additional recording burden exceeds its demonstrable benefit.

Formal representations do not themselves authenticate a record. A digitally signed event can certify who signed it under a security model while leaving the truth of the recorded ecological observation unresolved. Likewise, an inferred causal path can be useful for hypothesis generation without identifying the policy effect in Equation (75). Reviews should independently audit the completeness of the relevant causal graph, provenance, and material account; distinguish graph reachability from physical necessity; and report where a chosen generative relation or representation changes the conclusions. The appropriate comparison may show that familiar methods already answer a question more transparently. Such an outcome narrows the GR research programme in a scientifically productive manner.

## 27.8 Conditions for Rejecting the Proposed Comparisons

A useful policy study identifies, in advance, the observations that would count against its motivating explanations. If throughput optimization increases direct resource use at fixed output, the result may be attributable to run scheduling or idle-time management rather than to autonomous AI specifically; if energy optimization reduces lifecycle impacts as well as direct consumption, a hypothesized displaced upstream burden may be absent. If the ecologically constrained arm cannot produce the stipulated service, the result is infeasibility under the chosen technology, limits, and horizon, not proof that ecological protection is impossible in general. If no significant or decision-relevant effects are identifiable because uncertainty is large, the published conclusion should be indeterminacy with an explicit interval and study boundary.

The major failure conditions can be stated before data collection. First, a claimed policy effect fails identification when treated and comparison units lack overlapping eligible conditions, assignment or adherence is severely confounded, or interference is important but unmodeled. Second, a claimed ecological improvement fails when absolute burdens increase within a defensible extended boundary, a material ecological category exceeds its declared limit, or measurement quality does not support the relevant direction of change. Third, a claimed gain in autonomy fails when the policy relies on unrecorded human interventions or increases critical external maintenance or procurement dependence beyond the predeclared tolerance. Fourth, a claimed institutional explanation fails when changes in financing, access, or decision rights do not explain observable choices after plausible alternatives and selection are examined. Fifth, the claim of added methodological value fails when GR adds no reliable predictive, diagnostic, or provenance advantage over informed conventional methods at comparable informational and computational cost. These possibilities should be registered separately; they are neither logically equivalent nor necessarily observed together.

It remains essential to distinguish evidence for choosing *among feasible options* from authority to de-

DRAFT

cide *which outcomes deserve protection*. Preassigned ecological limits, employment conditions, community representation, and rights to remedy depend on scientific argument, legal institutions, and contested normative judgments. No scalar optimizer or causal effect estimate can establish that an unauthorized burden on another community is legitimate. Conversely, a finding that an ecological policy is expensive at one site is not sufficient to establish that its constraint is unnecessary. The comparative inquiry supplies explicit material, causal, and institutional evidence for deliberation, while preserving the distinction between descriptive, explanatory, and normative conclusions.

The result of this section is a testable research design rather than an empirical verdict. It specifies comparable service, executable treatment arms, credible baselines, causal targets, lifecycle boundaries, interference scenarios, governance conditions, uncertainty records, and potential negative findings. The next section addresses a remaining prerequisite: whether productive dependencies and ecological exposures can be measured with sufficiently reliable provenance in concrete automated laboratories, industrial cells, AI-compute supply chains, and maintenance networks to implement any of these comparisons. Until such case evidence is assembled, the present claims remain conditional protocols and hypotheses, not verified effects of autonomous production.

# 28 Measuring Productive Dependencies in Empirical Cases

Section 27 specified counterfactual comparisons among production policies with matched service requirements, explicit lifecycle boundaries, and prospective failure criteria. Such comparisons depend on a more elementary empirical achievement: identifying which resources, operations, services, authorities, and ecological processes the production system actually relies upon, and determining which descriptions can be supported by traceable observations. This section develops a modular measurement programme for five settings: an automated materials laboratory, an industrial production cell, an AI-compute and equipment supply chain, a maintenance network, and a production-linked ecological supply chain. It translates the typed relations of Section 6 into operational indicators, source-quality rules, and reproducible data products. Its aim is to establish conditions for collecting and evaluating evidence, rather than to claim that the proposed measurements have already been obtained. Published laboratory demonstrations and international statistics are useful anchors, but neither replaces event-level data, site-specific ecological observations, nor the institutional records needed for inference about reproductive autonomy.

## 28.1 Defining the Unit of Observation and the Dependency Claim

An investigator should declare the unit of observation before constructing a dependency graph. For an industrial cell it may be a production order or qualified cycle; for a self-driving laboratory it may be a target–experiment–characterization episode; for maintenance it may be a work order linked to an asset and failure mode. For upstream equipment procurement it may be a delivered part or supplier batch, while an ecological study can require a georeferenced sampling unit and interval. The time coordinate records the ordering and duration available from a particular clock and observation protocol, not a presumed universal historical metric. At the same time, comparisons across modules require stable definitions of the supplied service, actor, equipment, material grade, and quality outcome. The choice between a cycle, a kilogram of product, a qualified research finding, or a facility-year is substantive: it changes which failures, human interventions, and external inputs enter the denominator.

The minimal evidentiary object is a claim about a typed relation, accompanied by the observation or document that warrants it. Let $v \in \mathsf{V}$ denote a versioned object or activity record, $r \in \mathsf{R}$ a candidate relation, and $e \in \mathsf{E}$ an evidence item. Define a proposed record schema

$$e = \big(\mathrm{id}, \mathrm{source}, \mathrm{method}, \mathrm{unit}, I, \mathrm{place}, \mathrm{version}, \mathrm{uncertainty}, \mathrm{access}, \mathrm{claim}\big). \tag{78}$$

Here $I$ is the observation interval with timezone and resolution, place is the site or geography at the precision permitted by consent and security, and access distinguishes public, restricted, and non-disclosable evidence. The claim field identifies the asserted relation and whether it is observed, directly measured,

document-derived, estimated from a process model, transferred from a regional database, or hypothesized. The other fields have no universally interchangeable measurement unit: they are typed metadata governing how a later physical quantity or institutional assertion may be used. The description is compatible in spirit with the entity–activity–agent distinctions of the W3C PROV ontology and the FAIR stewardship principles, without presuming that a provenance graph establishes the underlying event's truth [196, 191].

Three claims require separate tests. *Observed input use* can be warranted by an appropriate flow meter, inventory movement, and timestamped production event. *Functional dependence* requires evidence that absence, delay, or degradation of an input changes the ability to supply the declared service, after allowance for buffers and substitutes. *Causal necessity* makes a stronger counterfactual claim about what would have happened if a particular condition had been altered; it generally requires an intervention, a defensible structural model, or a natural experiment with stated identification assumptions. A purchase order alone supports none of these stronger claims. Conversely, an unrecorded part can still be materially necessary. The graph of recorded relations is therefore a map of supported claims, not a complete inventory of physical reality.

A dependency audit should specify conjunctions, alternatives, and thresholds. A robot cell may require both qualified power and a compatible tool, while replacement tools may be obtained from either of two suppliers. An upstream dependency should not be counted as an immediate production halt where validated substitute inputs and inventories provide continuity over the declared horizon. For a candidate input $r$, a proposed *model-based service-loss contrast* is

$$\Delta_r(\pi, H \mid \mathcal{M}) = \Pr_{\mathcal{M}}[Q_H < q^* \mid \mathrm{do}(A_r = 0), \pi] - \Pr_{\mathcal{M}}[Q_H < q^* \mid \mathrm{do}(A_r = 1), \pi]\,, \tag{79}$$

where $Q_H$ is conforming service in declared functional units over horizon $H$, $q^*$ is the required quantity, $A_r$ indicates availability of the specific input or enabling relation, $\pi$ is the operating and substitution policy, and $\mathcal{M}$ is an explicitly stated intervention model. The expression is dimensionless, conditional on $\mathcal{M}$, and need not be positive in a complex system with responses and substitutions. Neither a graph path nor an empirical correlation identifies $\Delta_r$. Sensitivity to repair arrangements, reserve capacity, and supplier substitution should be reported rather than compressed into an unsupported single criticality score.

## 28.2 Reconciling Physical Flows with Activity Records

Physical and event-based records should be joined through shared, audited identifiers rather than by coincident timestamps alone. A proposed minimal industrial data package contains: asset and sensor registries; calibration certificates; order and batch specifications; controller and software versions; robot and machine-state events; material receipts, withdrawals, scrap, and shipments; time-resolved electricity and fuel; water and process chemicals where relevant; quality inspections; operator interventions; maintenance events; supplier lot genealogy; and explicit boundary registers. IEC 62264-1 provides enterprise–control interface models, ISO 22400-2 manufacturing performance-indicator definitions, and ISO 23247 a manufacturing digital-twin framework [62, 86, 89]. NIST's completed Digital Thread for Manufacturing project provides research context for the integration of design, manufacturing, and inspection information rather than an additional normative manufacturing data standard [133]. Neither the standards nor the project description establish that the required variables are populated, correct, or independently reproducible in a particular installation.

A useful first check is the conservation residual. For material species or grade $k$ recorded in kilograms, let $x_{k,0}$ and $x_{k,H}$ be measured inventory, $m_k^{\mathrm{in}}$ measured receipts or generation, and $m_k^{\mathrm{out}}$ measured shipments, releases, transformation-consistent consumption, or other documented sinks over the same horizon. Write

$$\varepsilon_{k,H} = x_{k,H} - x_{k,0} - m_k^{\mathrm{in}} + m_k^{\mathrm{out}} \qquad [\mathrm{kg}], \tag{80}$$

with a suitable inventory definition and a consistent account of transformations between materials. A nonzero residual may indicate timing differences, unregistered flows, calibration bias, grade mismatch,

DRAFT

Table 49: Proposed observations for the empirical modules. Public sources can motivate a case but do not substitute for the new primary measurements shown here. FU denotes a study-specific qualified functional unit.

| Module and unit | Observable dependency question | Primary records still required | Available public context |
|---|---|---|---|
| Autonomous materials laboratory; validated experiment | Which manual, computational, reagent, characterization, and maintenance services enable each accepted finding? | Run and sample provenance, human interventions, reagent lots (kg or mol), kWh, failures, instrument calibration | Published A-Lab demonstration and correction; laboratory standards |
| Industrial production cell; conforming FU | Do software, tools, spares, supplier conditions, and electricity limit qualified production across failures? | Timestamped process and quality logs, metered energy (kWh), scrap (kg), orders, downtime, worker interventions (h) | IFR aggregate deployments; manufacturing KPI and interface standards |
| AI-compute and equipment chain; defined compute service | Which accelerator, memory, grid, cooling, supplier, and replacement dependencies are concentrated or substitutable? | Facility meters, compute workloads, equipment bills, serial-level repair, embodied inventories, procurement records | IEA scenarios; public trade and equipment information |
| Maintenance network; restored asset-service hours | How do skills, parts, diagnostics, contracts, and repair lead times determine recovery? | Asset genealogy, work orders, failure causes, lead time (h), parts, labor, restoration tests | Reliability and maintenance standards; public benchmark datasets |
| Ecological supply chain; qualified FU and receptor | How do extraction, production, transport, water use, and discharge intersect identifiable ecological receptors? | Geocoded facility and batch links, flow inventories, exposure and ecological monitoring, baseline and lag records | UNEP resource accounts; environmental databases and local registers |

or genuine measurement error. It cannot by itself identify an ecological release or dishonest record. Analogous energy balances require properly distinguished electricity, heat, fuel, conversion, and storage terms, with compatible energy units. Investigators should retain the raw readings and their uncertainties, calibration dates, gap-filling rules, and reconciliation decisions. For equipment mass and embedded energy, physical flow observations may need process-based lifecycle inventories whose boundaries and allocation choices remain separately visible under ISO 14040/14044 [84, 83].

Table 49 makes two points that bear upon feasibility. First, *directly observable* means accessible through an appropriate instrument or record, not necessarily public. Second, a public macroeconomic or technology statistic cannot be assigned without qualification to a particular facility. Location, observation date, technology, and sampling convention are attributes of the evidence, not decorative citations. A model relying on regional means should mark them as background estimates and show how estimates change under plausible site conditions. The best outcome of an initial audit may be a narrower, reproducible system boundary rather than an apparently comprehensive but unverifiable network.

## 28.3 An Automated Materials Laboratory as a Bounded Research Case

The autonomous inorganic-materials laboratory reported by Szymanski and colleagues provides a documented starting point because it integrated literature-derived synthesis planning, calculations, active learning, robotic processing, and characterization. The *corrected* published article, following an author correction issued in January 2026, reports 36 compounds realized among 57 targets over 17 days of operation, rather than the earlier 41 among 58 figures [178, 179]. Those are author-reported results for the particular experimental system and its success criteria; they are not observations of end-to-end mineral sourcing, electricity infrastructure reproduction, laboratory-independent scientific discovery, or ecological compatibility. Equally, a synthesis outcome and a fully validated new materials claim may require different analytical criteria. This case is suited to designing, rather than claiming completion of, an extended measurement record.

DRAFT

A proposed laboratory audit would begin at the eligible target and preserve the genealogy of every executed experimental episode: computational prediction and its dataset version; proposed recipe; requested reagent amounts and grade; human approval or editing; robot task execution; consumed crucibles or disposable vessels; furnace heating and cooling; instrument characterization; interpretation; rejection or repetition; and the final independent quality determination. The register would distinguish fully automated tasks from human preparation, troubleshooting, calibration, hazardous-material handling, safety supervision, interpretive correction, and waste treatment. For scientific service, a count of experiments may be reported alongside a prespecified validity rubric, rather than substituted for scientifically meaningful output. Reviews of self-driving laboratories and proposed performance metrics motivate task- and study-specific autonomy measures [20, 188]. ISO/IEC 17025 separately specifies requirements for laboratory competence, impartiality, and consistent operation [96]; citing it does not certify any experiment or establish that a particular laboratory satisfies those requirements.

The laboratory's energy inventory would separate compute, instruments, heating, ventilation, standby, and shared building loads where metering allows; allocation of shared loads would be marked with a model rather than invented measurements. Material accounting would distinguish incorporated compounds from rejected products, consumables, reagent packaging, waste treatment, and replacement instrumentation. A run recorded as autonomous should still show whether personnel replenished powders, moved samples, corrected model outputs, or resolved failed instruments. Public documentation alone rarely supports complete reconstruction of these histories. Direct access to original logs and supplier information would be necessary; where intellectual-property, safety, and laboratory-privacy limits prevent disclosure, a secured audit or withheld-data declaration is preferable to attributing fabricated completeness to an openly available paper.

A practical negative result would be the inability to link reported outputs to a sufficiently complete genealogy of labor, energy, consumables, and instrumentation to estimate reproductive autonomy on a declared horizon. Another possibility is that independent scientific verification changes the accepted-output denominator. Both affect interpretation of productivity and scientific service, while leaving the published local automation achievement as a distinct, bounded claim. The proposed study therefore retains platform capability, knowledge reliability, and material reproduction as three independently evidenced dimensions.

## 28.4 An Industrial Cell and the Difference between Operating and Reproductive Continuity

An industrial-cell case can combine an articulated robot, a numerically controlled processing machine, automated inspection, and a material-handling subsystem. This configuration is a *proposed study design*, not a claim about a presently accessible factory. A reasonable comparator is the same cell operated with a conventional, competent control and maintenance regime. The investigators predefine eligible products, tolerances, quality inspections, safety controls, and operating intervals, then trace both nominal operations and specified disturbances. Relevant disturbances include missing replacement parts, calibration drift, supplier delay, loss of communications, coolant interruption, software-licence or support constraints, and unavailable qualified staff. These are hypothesized stress scenarios and require site-specific safety review before any actual intervention.

The International Federation of Robotics reported in its September 2026 *World Robotics 2026* release that worldwide factory installations exceeded 600,000 robots during 2025 and the operating stock reached approximately five million at year-end [69, 70]. These are global aggregated installation and stock measures with different denominators and are not samples of event-level downtime, material use, controller autonomy, or ecology in the proposed industrial cell. They substantiate the contemporary relevance of instrumentation and comparison but do not reveal its results. Industrial-cell evidence instead requires production timestamps, independently measured conforming output, component and tool genealogy, metered energy at specified resolution, machine downtime, maintenance hours, and quality records. Robot performance characteristics, robot-design safety requirements, and manufacturing-operations KPIs have

DRAFT

distinct standards bases [81, 94, 86]. These standards do not measure the employment quality or effective decision rights of workers, which require separate observation and instruments. Moreover, ISO 10218-1:2025 addresses the industrial robot itself; the safety of an integrated cell requires its own application-level assessment.

One analytic target is *service continuity across a disturbance*. Define $T_{\text{loss}}$ as the total hours in which the cell fails its specified required service rate within horizon $H$, together with the qualified units not delivered. Record fault occurrence, diagnostic delay, parts procurement, hands-on repair, requalification, and restart as separate intervals. The service record can then be compared against the presence of trained workers, spares, validated alternative suppliers, remote support, and internal maintenance authority. This does not equate any observed association with the causal effect in Equation (79); it provides variables and events for a later design capable of identifying that effect. An ostensibly self-operating cell may exhibit high nominal uptime while relying heavily on an externally controlled precision component. Conversely, automation may facilitate local diagnosis and rapid recovery. The measurement programme admits either finding and records changes in skill requirements and authority as well as equipment availability.

## 28.5 The Distributed Material Reproduction of AI Computing

The third module extends the observation boundary beyond the visible production cell. AI-assisted production may rely on training or inference services, accelerators, memory, semiconductor fabrication, electricity, transformers, cooling, networks, and specialized equipment repair. A defensible unit of service is a defined class of validated inference or planning jobs at specified accuracy and latency, paired with an observation period for provisioning and replacement capacity. Raw floating-point operation counts alone may not capture the relevant quality of the supplied production decision. Moreover, the source of power and water, the provenance of compute hardware, and the capacity to replace failed equipment may differ across locations and years. The methodological challenge is to join these distinct layers without pretending that they are all directly measured by the same operator.

The International Energy Agency's 2026 report *Key Questions on Energy and AI* provides an updated, explicitly modeled outlook of data-centre electricity requirements and associated bottlenecks; the earlier *Energy and AI* analysis identifies its estimates and projections as uncertain and scenario dependent [68]. An IEA global estimate cannot be interpreted as metered electricity for a selected inference workload or manufacturing facility. Facility instrumentation would need to record job-level allocation where feasible, IT power, shared cooling and conversion losses, equipment utilization, grid context, and any contractual renewable-power attribution. A cradle-to-grave inventory would then include equipment manufacturing and replacement, semiconductor and critical-mineral supply, transport, and end-of-life outcomes, with background processes drawn from versioned LCA and trade datasets. Public critical-minerals assessments describe risks at sector scale but cannot disclose the actual supplier path of a particular accelerator without additional documentation [67, 141, 83].

The data architecture should therefore distinguish *facility-measured* electricity (kWh), *workload-allocated* energy (kWh per defined service under a declared allocation model), *background-estimated* embodied burdens (units and characterization method specified), and *scenario-modeled* marginal consequences. Supplier concentration requires a supplier-tier register with geography, product compatibility, available alternatives, realistic qualification delays, and a clearly stated horizon. A sector-level Herfindahl index or national export concentration is insufficient to characterize the recoverability of a specific machine when procurement restrictions, maintenance licenses, inventories, and qualification differ. Undisclosed subcontractors and sensitive procurement terms should be coded as unobserved or restricted rather than as absent dependencies. Assessing power through a procurement graph alone would additionally omit institutional permission to alter software, access diagnostics, obtain spare parts, or move work to a substitute provider.

The module's primary contribution would be auditable *bridges* among physical operation, embodied inputs, supplier-specific constraint scenarios, and institutional control. A well-instrumented facility might show that external compute is negligible for a particular task, or that the largest burden lies in non-

DRAFT

AI process heating. Either conclusion is permissible. The purpose is not to reassign the entire global computing footprint to an individual production activity; it is to establish the smallest defensible account and identify the upper tiers for which only modeled sensitivity ranges can presently be supplied.

## 28.6 Maintenance Events and the Reproduction of Repair Capacity

The fourth module investigates a comparatively neglected component of productive autonomy: how damaged, depleted, or unreliable equipment is restored to qualified service. Maintenance studies already have a substantial literature on condition monitoring, reliability, diagnostic accuracy, asset management, and remaining useful life [104, 110, 87, 92]. The proposed relational account adds an explicit attempt to observe how repairs depend on parts inventories, external technical documentation, personnel skills, contracting, and rights to access diagnostics. This should be evaluated against an appropriately detailed computerized maintenance management system and reliability model, rather than against a repair ledger deprived of procurement information.

The observational unit is a work order linked to the affected asset, a predeclared failure category, the production service impaired, and a validated restoration event. A minimum record includes the preceding duty cycle or equipment age; detection method; detected and independently adjudicated fault; time from detection to authorized intervention; personnel skill and hours; diagnostic software and access permission; spares requested and consumed; supplier lead time; repair and requalification time; and recurrence or rework within a defined follow-up. The outcome is not merely a closed ticket. A repair is successful when the asset meets an independently specified operational and safety acceptance criterion; an interim bypass or incomplete corrective action should be tagged separately. Distinguish scheduled maintenance, corrective repair, emergency workaround, and replacement, since mixing these interventions can obscure whether an apparently low downtime reflects effective prevention, a favorable asset population, or deferred maintenance.

For asset $i$ and observation horizon $H$ measured in hours, a basic, noncausal continuity descriptor is

$$\widehat{A}_{i,H} = \frac{H - T_{i,H}^{\text{unavailable}}}{H}, \qquad 0 \le \widehat{A}_{i,H} \le 1, \tag{81}$$

where $T_{i,H}^{\text{unavailable}}$ is the measured unavailable time under a declared scheduling convention, excluding and including planned outages in separately reported versions where useful. This is a descriptive availability fraction, not a causal measure of autonomy, and it can differ from the fraction of qualified production delivered when buffers or redundant assets compensate for an outage. An operator's ability to restore equipment without vendor assistance should be recorded as a distinct, capability-specific question. Failure logs often contain selected cases only, and maintenance intensity is itself a response to equipment condition. A naive comparison of repaired assets to unrepaired assets may therefore invert the apparent effect of maintenance; selection and time-varying treatment require the design discipline developed in Section 27.

A research protocol can stratify by asset family, duty cycle, workload, failure severity, and contract status. It should track repairs completed with local parts and expertise, repairs requiring supplier delivery, remote diagnostics requiring a proprietary credential, and cases with no feasible repair within the observation horizon. Component lead-time distributions and authorized substitutions may be more informative than an average repair duration. A narrowly defined hypothetical example illustrates the distinction: suppose an identical machine has two qualified suppliers for a consumable part and a single proprietary diagnostic service. The presence of two part suppliers may eliminate a particular material bottleneck, while an unavailable diagnostic credential still prevents restoration. This logical example supplies no frequency, likelihood, or empirical evidence; it demonstrates why availability is a conjunction of material, knowledge, skill, and authority conditions. The empirical test would be whether actual event records and independently checked compatibility support this structure.

DRAFT

## 28.7 Connecting Production Histories to Ecological Receptors

The fifth module asks whether upstream procurement and local production records can be linked to a defensible ecological observation boundary. Its case is a *prospective site-specific study* involving a traceable material batch, its processing and transport, resource and water use, emissions, waste handling, and one or more potentially exposed ecosystems. A material-flow account can establish physical throughput and the documented location of a discharge. Characterized life-cycle impacts can approximate pressures under a selected model. Observed biotic or hydrological changes require a different evidentiary apparatus, with appropriate baseline, comparison, temporal lags, spatial exposure pathways, and relevant ecological endpoints. The International Resource Panel's *Global Resources Outlook 2024* provides a global resource-use context; its aggregates do not identify localized ecological damage or recovery from a particular industrial intervention [184, 59].

The proposal separates three spatial levels. A site-level inventory specifies operations, receiving water body or other receptor, withdrawal and discharge locations, and monitoring sites in the appropriate coordinate reference system. A supply-chain extension uses documented origins and transportation links, retaining unknown provenance as missing rather than applying unqualified average locations. A regional or global extension draws on background lifecycle or multiregional input–output databases with acknowledged resolution limits [173, 190]. For water withdrawal and consumption, seasonal and catchment context may change interpretation, while a greenhouse-gas inventory generally addresses a different spatial mechanism; ISO 14046 supplies relevant assessment conventions for water footprints [85, 83]. A remote supplier's reported carbon intensity cannot reveal a local loss of habitat connectivity or the condition of the receiving ecosystem.

Ecological field observations should be selected around a mechanism, such as a specified contaminant's transport pathway, a water-flow change, or a habitat disturbance linked to land conversion. Before sampling, investigators declare affected and plausible comparison sites; measurement precision and observation frequency; background variation; relevant seasons; ecological lag windows; and criteria for possible recovery. Consent and consultation with affected communities, scientific sampling permits, the protection of vulnerable species locations, and any legal access restrictions are methodological conditions for responsible data collection. A field study cannot assign a causal effect merely because a biological decline follows an industrial expansion. Statistical controls, intervention timing, upstream trends, multiple stressors, and appropriate ecologically meaningful replication would all require scrutiny. A meta-analysis of restored wetlands finds long-lasting shortfalls in biological structure and biogeochemical functioning relative to reference sites [129]. A separate synthesis of tropical-forest studies compares natural regeneration with active restoration using biodiversity and vegetation-structure measures [25]. Neither result supplies a universal recovery timescale for an industrially affected site; short favourable indicators should not be equated with full restoration.

To prevent false commensuration, the resulting dataset should retain a vector of outcomes: physical inflows and releases by substance and unit, characterized impacts by declared method, directly observed exposure, and ecological indicators with their site and sampling definitions. The relation between these classes may be modeled, but the model must be separately evaluated and must not treat the ecological indicator as a direct reading from the factory's energy meter. Ecological thresholds involve scientific uncertainty and normative or legal choices about acceptable loss. The research task is to expose these choices, test the support for their material premises, and preserve contrary or incomplete evidence.

## 28.8 Source Provenance, Missingness, and Bounded Coverage

The five modules require a common audit procedure because missing data are patterned by institutional access as well as technical measurement capacity. Supplier inventories may be commercially confidential, operator interventions may go unlogged when they are informal, upstream ecological conditions may never have been sampled, and an industrial site may retain high-frequency operating data while deleting its older software versions. These gaps should be classified by source type and mechanism: not collected, inaccessible, lost, incompatible, below detection limit, estimated, imputed, or structurally

DRAFT

Table 50: Evidence classes, permissions, and legitimate uses in a dependency audit. The proposed classification is an operational coding scheme, not an established statistical scale.

| Evidence class | Example and audit requirement | Permitted initial use | Inference withheld without additional evidence |
|---|---|---|---|
| Direct measurement | Calibrated meter or adjudicated test; interval, unit, uncertainty | Local measured flow or service with known boundary | Upstream completeness or causal necessity |
| Event/document observation | Time-stamped work order, verified contract, licence, dispatch log | Descriptive sequence or documented authority | Actual delivery, lawful enforceability in every jurisdiction, causal effect |
| Site-specific estimate | Explicit allocation of shared cooling or inferred material balance | Conditional inventory with assumptions and sensitivity | Independent measured value |
| Transferred background estimate | Regional electricity mix, IO table, LCA unit process | Modeled external extension with database version and coverage | Supplier- or batch-specific observation |
| Scenario/structural assumption | Hypothetical substitution, disturbance or exposure response | Counterfactual modelling and sensitivity tests | Observed outcome or validated causal pathway |
| Missing/restricted | Unavailable supplier tier or unshared calibration certificate | Declared scope limitation and possible bounded analysis | Verified absence of a dependency |

inapplicable. There is no reason to expect the missingness mechanism to be independent of the presence of a critical dependency. A model trained on facilities with permissive data access may consequently generalize poorly to firms or regions where disclosure incentives differ.

Even a rigorous provenance register cannot certify that its universe of relevant dependencies is complete. A transparent *registered-field coverage* measure can be defined without claiming to know the total unseen network. Let $\mathcal{K}_B$ be a prespecified set of $K_B > 0$ eligible indicator fields within declared boundary $B$ (for example the required entries for all registered work orders), and let $m_j = 1$ when field $j$ has an admissible, independently auditable value at the necessary resolution, zero otherwise. Then

$$C_B^{\text{reg}} = \frac{1}{K_B} \sum_{j \in \mathcal{K}_B} m_j, \qquad 0 \le C_B^{\text{reg}} \le 1. \tag{82}$$

$C_B^{\text{reg}}$ is dimensionless and descriptive. Its denominator refers to the *registered* data requirements, not to all actual physical, institutional, or ecological relations. A score of one is compatible with substantial unidentified supplier tiers if those tiers were omitted from the registry. For this reason the audit reports separate registered-field coverage, population coverage (for example the fraction of eligible assets or orders represented), linked-flow mass/energy reconciliation, and open-boundary flags. The investigator should periodically challenge the eligibility register using independent sampling, procurement records, site visits where permitted, supplier questionnaires, and comparison against competent industrial-ecology baselines. Improvements that derive solely from redefining an unmeasured field as inapplicable should be explicitly disclosed.

LCA standards require attention to goal, scope, inventory, interpretation, and study limitations [83]. Practical data-quality assessment may further describe reliability, completeness, and temporal, geographical, and technological correspondence, with criteria defined by the chosen methodological guidance [183]. A particular scoring scheme should not be attributed to the title or abstract of ISO 14044 alone. By contrast, JCGM 100:2008 concerns the evaluation and expression of measurement uncertainty; its quantitative treatment should be applied to defined measurands and uncertainty models, not treated as

DRAFT

Table 51: Staged empirical outputs and failure conditions. Each stage has an independent acceptance criterion; progression does not imply that its claims are already validated.

| Gate | Deliverable | Minimum evidence for proceeding | Failure or qualified result |
|---|---|---|---|
| 1. Boundary | Functional unit, asset/site roster, recorded service window | Reproducible unit definitions and source access statement | Incommensurate services; unresolved case selection |
| 2. Inventory | Validated material/energy/maintenance event extracts | Calibration and reconciliation with missingness ledger | Unexplained residuals; unsuitable resolution |
| 3. Relation | Adjudicated typed dependence and substitute paths | Evidence class and source for each asserted consequential relation | Documentation mistaken for functional necessity |
| 4. Extension | Versioned upstream and ecological boundary scenarios | Representative source, geographic/time match, uncertainty ranges | Background data wrongly reported as site measurement |
| 5. Comparison | Preregistered conventional-model and GR-augmented analyses | Blind evaluation on held-out histories with equal information | No GR incremental benefit; leakage or unjustified overhead |

a validated score for documentary truth [107]. Both sources of uncertainty should inform sensitivity analysis rather than become a fictitious universal probability that a recorded statement is true. Where multiple sources refer to one event, investigators should document conflict resolution: for example, a supplier invoice, machine log, and operator interview may record distinct dates without one automatically overruling the others. Recorded timestamps should preserve source clock, timezone, precision, and synchronization assumptions. Every transformed estimate should reference its source data, code version, method, reviewer decision, and permitted reuse. When proprietary or personal information cannot be released, a data dictionary, non-sensitive aggregates, controlled verification, and a reproducible synthetic example are more defensible than claiming public auditability of withheld records.

## 28.9 A Reproducible Study Package and an Incremental Representation Test

The immediate output of the proposed programme is a versioned *study package* rather than a single indicator. At a minimum it should contain a case selection statement; boundaries, functional units, data-access agreements, permissions and risks; a typed entity and relation dictionary; a sampling and calibration plan; event and flow schemas; raw-to-analysis transformation specifications; a missingness register; process and ecological measurement methods; prespecified comparison outcomes; and a ledger of revisions. Data provenance can follow interoperable W3C concepts, but an independently generated evidence identifier is needed to prevent two duplicated claims from being counted as corroboration. Public code should include unit tests for dimensional compatibility, impossible timestamps, duplicated identifiers, unmatched input/output events, and disagreement between equipment state and recorded service. Where raw data are restricted, independent evaluators need a defined process for verifying the released claims within approved access arrangements. PROV-O addresses provenance relations and FAIR addresses discoverability and reuse; NIST's completed digital-thread project documents a specific manufacturing-information integration problem [196, 191, 133]. None of these resources authorizes disclosure of restricted data, independently authenticates a recorded event, or supplies an industrial causal identification strategy.

The final gate provides a test of whether the proposed GR representation supplies information beyond mature alternatives. Baseline A should comprise a well-implemented bill-of-materials and asset registry

DRAFT

linked to process measurements, LCA/MFA, supply-chain mapping, and maintenance analysis. Baseline B should add a conventional temporally indexed production network or a manufacturing digital twin developed using published framework principles, with provenance metadata and relevant authorization and maintenance fields [88, 196]. The experimental candidate C can then add the typed multi-party activity, authorization, ecological-receptor, and event-history features proposed in Sections 6 and 26. All alternatives must receive equivalent source data, and baseline investigators should be allowed to engineer strong features. Evaluation tasks include recovering the documented sequence of a material batch, identifying verified dependency failures, predicting service interruption in held-out disruptions, localizing unaccounted emissions within a known boundary, and diagnosing whether an institution could execute a stated ecological constraint. Ground-truth labels should be independently adjudicated, with uncertain cases excluded from accuracy scoring or scored according to a declared uncertainty protocol.

Incremental performance should be measured by task-specific precision, recall, predictive error, calibration, correct identification of alternative suppliers, provenance-completeness diagnostics, analyst time, and computational as well as data-collection cost. Publication of favorable case studies alone is insufficient. The experimental candidate fails to establish incremental merit when it cannot outperform the competent baselines, when a gain disappears after equivalent provenance and authority data are supplied to the baselines, when gains are confined to mislabeled training/test duplicates, or when the maintenance cost of the model outweighs a demonstrable analytic benefit. This is a comparison of *representations used for stated tasks*, not a metaphysical experiment about whether the production world is fundamentally generative-relational. A prespecified ablation should evaluate whether any advantage comes from typed conjunctions and qualified alternatives, temporal/event histories, institutional authorization links, or ecological-receptor mappings; dropping each component from the candidate and adding the corresponding information to the baseline must be evaluated on the same held-out cases. A result that depends on unequal records, uncontrolled analyst effort, or favorable retrospectively selected tasks would not meet this gate. Three adversarial cases should be registered before implementation. In the first, a competent baseline receives identical provenance, authorization, and repair-event fields through an extensible ordinary database: any purported representational advantage must survive this information-equivalent extension. In the second, the same supplier interruption occurs in several highly related facilities: train–test partitioning must keep the supplier-event cluster intact, or repeated records may masquerade as independent predictions. In the third, a GR workflow identifies more environmental fields only because it funds additional sampling: report its expanded coverage and acquisition costs separately from the performance comparison on the shared observation set. A fourth check should recompute results after changing a defensible boundary or terminal horizon; a claim that reverses requires scope-qualified reporting. These are designed stress tests, not outcomes of a conducted experiment.

For a first preregistered benchmark, three independently adjudicated tasks are preferable to an unspecified composite score: (i) prospective interruption prediction after independently logged part or service failures, scored for service-loss error and predictive calibration; (ii) reconstruction of a documented material or maintenance lineage, scored for verified relation precision and recall with explicit unknown cases; and (iii) assessment of whether a stated production-policy alternative was actually exercisable, scored against contemporaneous contractual and operational evidence rather than a retrospective analyst opinion. The third task must not turn an institutional judgment into an automated declaration of legitimacy. A competent digital-twin or temporal-graph implementation should be allowed to represent qualified substitutes, authorization fields, and uncertainty, rather than being limited to a deliberately impoverished graph. Where the ecological module is included, a separate physical and lifecycle-accounting assessment remains necessary; a provenance benchmark is not a direct test of ecosystem effects.

The protocol should register the minimal meaningful difference separately for each outcome, the number and grouping of held-out facilities and disturbance episodes, the handling of dependent events from one facility, the planned uncertainty intervals, and a resource ceiling for both approaches. Training and configuration must not use future revisions of the adjudication records. Findings should be reported by task and site, with negative and inconclusive results retained, rather than selectively pooled into a single overall success percentage. A model that yields equal accuracy but better independently assessed

DRAFT

contestability may have a *procedural* contribution; one that yields more accurate predictions only by using superior source access has not demonstrated a *representational* contribution. Claims of reduced ecological burdens require their own independently measured or causally justified evidence.

### 28.10 What the Empirical Programme Can Establish

The modular suite yields three types of possible finding. First, direct observations can establish bounded inventories and sequences: electricity consumed at a measured facility, tools and components exchanged, operator hours on specific work orders, documented approvals, and changes in recorded quality and service. Second, validated engineering models and appropriate comparison designs can estimate how alternative inputs, repairs, and policies influence service continuity, provided that the inferential assumptions survive testing. Third, lifecycle and field-ecology methods may relate physical burdens to wider supply chains and environmental receptors, with explicit geography, time horizon, uncertainty, and normative boundary choices. These levels of claim remain distinct even when their evidence can be joined through a typed representation.

A successful inventory does not alone demonstrate the self-reproduction of production infrastructure, and a plausible ecological exposure path does not alone establish its causal impact. Published aggregate statistics, including robot stock and international energy demand, can define context and external plausibility limits, but do not authorize reconstruction of individual organizations' dependency histories. The empirically serious question is whether the declared instruments, public databases, institutional permissions, and ecological observations permit meaningful tests of proposed differences between operational autonomy, infrastructural dependence, and ecologically viable productive continuation. Absence of sufficient evidence is itself an outcome to be reported with the unresolved fields, assumptions, and decision consequences made legible.

This section thus supplies an operational bridge between the counterfactual design of Section 27 and the critical interpretation required for the paper's conclusion. It defines auditable units of service, distinguishes relation observations from claims of necessity, specifies complementary physical and institutional records, identifies five technically differentiated study modules, and sets reproducible source-quality and comparison gates. Its demonstration is methodological: no new experimental dataset, field survey, causal estimate, or measured ecological improvement has been generated. The following section will assess the circumstances in which missing supplier tiers, selected facilities, unreliable ecological baselines, model structure, and institutional asymmetries make even a carefully organized comparison indeterminate.

## 29 The Limits of Inference in Comparative Production Research

Section 28 proposed an empirical programme for relating production events, material and energy accounts, supplier networks, repair histories, institutional decisions, and ecological observations. Even well-organized records leave a further difficulty: the events available for analysis are a selected, incomplete, and often institutionally mediated representation of the processes to which a scientific conclusion refers. This section examines the resulting limits of inference. Its purpose is to determine when observations warrant descriptive inventories, dependency claims, causal comparisons, and claims about ecological continuation, and when the research must instead report qualified bounds or unresolved alternatives. The inquiry addresses missing supplier tiers, the selection of unusually capable facilities, proprietary data, ecological baseline gaps, competing network models, and power-conditioned observability. It concludes by registering prospective disconfirmation and non-identification criteria for the paper's principal hypotheses. The exercise is methodological and critical: no original comparison dataset or estimated empirical effect is introduced.

DRAFT

## 29.1 Separating the Questions that Evidence Can Answer

An observed production event supports a statement about the event, subject to the reliability of its instrument and record. A joined inventory supports a statement about a defined flow account, conditional on classification, measurement, and boundary coverage. A dependency claim adds a hypothesis about how a productive service would change if an input or relation changed. A causal policy comparison adds an intervention and the unobserved potential outcome under a credible alternative. Finally, an evaluation of whether a production regime is ecologically acceptable or institutionally legitimate requires further judgments about impacts, rights, authority, and the distribution of burdens. These evidentiary steps have different identification requirements. A complete event log does not entail an identified causal effect; an identified facility-level effect does not entail an identified regional ecological outcome; and the latter would not alone determine a legitimate decision.

For clarity, let $\mathcal{O}$ be the collected evidence and its provenance, $\mathcal{A}$ the maintained measurement, sampling, model, and causal assumptions, and $\Theta$ a prespecified set of target quantities or propositions. The *identified set* of values compatible with the stated evidence and assumptions is written

$$\mathcal{I}(\mathcal{O}, \mathcal{A}; \psi) = \{\vartheta \in \Theta : \exists m \in \mathcal{M}(\mathcal{A}) \text{ consistent with } \mathcal{O} \text{ and satisfying } \psi(m) = \vartheta\}. \tag{83}$$

Here $\mathcal{M}(\mathcal{A})$ denotes the candidate-model class admitted by assumptions $\mathcal{A}$, and $\psi : \mathcal{M}(\mathcal{A}) \to \Theta$ is the fixed target functional mapping a model to the parameter value $\vartheta$. The consistency relation, model class, and target functional must be specified for a particular empirical question. Equation (83) is a definition, not an estimator or a proof that any proposed production model is correct. A singleton identified set permits point identification *under the assumptions*; multiple compatible values call for partial identification or explicit model dependence. An empty set can indicate contradictory records or assumptions, rather than the impossibility of the productive system itself. This approach follows the distinction between available data, assumptions, and identification emphasized in the statistical literature [118, 54].

The appropriate inferential output should accordingly be stated as one of: a directly documented finding within a declared boundary, a model-dependent estimate with uncertainty, an interval or set of compatible results, a counterfactual claim with defended identification conditions, or a proposition not presently testable. A *negative finding* differs from a *missing finding*: reliable observations that contradict a stated hypothesis supply negative evidence, whereas missing observations may leave it indeterminate. These categories should appear in the evidence ledger for every consequential conclusion. They also prevent the relational vocabulary of nodes, processes, and trajectories from being mistaken for an independently authenticated ontology of production.

## 29.2 How Records Become Missing and Why That Matters

Missing production data are heterogeneous. A supplier may never record a subcontractor's energy consumption; a machine sensor may be offline during a breakdown; a site may retain only successful experiments; maintenance work may be documented outside the main operating system; an ecological receptor may not be sampled before construction; and a governing organization may prohibit access to procurement or consent records. These mechanisms influence which parts of the network are visible. If records disappear precisely when facilities fail, treating the remaining histories as representative can understate dependence. If accessible suppliers are concentrated in jurisdictions with stringent reporting obligations, imputation from those suppliers may mischaracterize unobserved tiers. The concepts of missing completely at random, missing at random conditional on observed variables, and non-ignorable missingness supply useful distinctions, but the assumed mechanism requires substantive defense [114].

Let $R_{jt} = 1$ indicate that a prespecified relation or flow record $j$ at observational interval $t$ was captured in an analyzable form, and $R_{jt} = 0$ otherwise. Let $X_{jt}$ denote observed pre-capture attributes, $Y_{jt}$ a target quantity such as an input flow or interruption severity, and $Z_{jt}$ an unrecorded commercial, operational, or institutional condition. An illustrative observation process is

$$\Pr(R_{jt} = 1 \mid X_{jt}, Y_{jt}, Z_{jt}) = \rho(X_{jt}, Y_{jt}, Z_{jt}), \qquad 0 \le \rho \le 1. \tag{84}$$

DRAFT

A sample can support inference using observed covariates only under assumptions that make dependence on unavailable $Y_{jt}$ and $Z_{jt}$ ignorable, or by deploying methods that explicitly bound or model such dependence. A *recorded-edge density* is therefore only a coverage measure relative to a declared roster; it is not the fraction of all material dependencies actually present in the world. In particular, an unobserved link must be labeled *unknown*, not *absent*, unless an independently validated detection procedure establishes non-occurrence at the relevant sensitivity and resolution.

The study should distinguish *structural noncollection* from *temporary loss* and *restricted access*. Structural noncollection calls for changing the instrument or bounding the unknown. Temporary loss may allow calibration or repeated measurement; its association with failures must nevertheless be checked. Restricted access requires a documented permission regime and, where possible, independent disclosure or secure-audit procedures; withheld material should not be silently recoded as a zero input. For each field, the proposed audit records eligibility, collection instrument, collection interval, source owner, known nonresponse, access status, and whether the evidence was independently corroborated. Sensitivity scenarios can then posit plausible values for the missing elements and examine whether the inference changes. Multiple imputation by itself does not establish an ignorability assumption and must not be represented as recovering unseen suppliers from nothing.

## 29.3 The Selection of Facilities and the Transport of Findings

Automated laboratories and advanced robot cells described in scholarly demonstrations commonly have exceptional technical staff, specialized funding, calibrated instruments, or unusually stable infrastructure. A demonstration can establish that a task was performed in a particular setting; it need not establish that the same productive sequence will operate in a geographically distant factory with different wages, power reliability, replacement parts, environmental conditions, or institutional support. The sampling frame matters as much as experimental sophistication. A transport formula cannot rescue a target population whose key production states, maintenance regimes, or ecological exposures fall outside the studied covariate support. Positive publicity and publication practices can also select conspicuous successes over interrupted or unsuccessful deployments. Research on external validity addresses precisely why an internally credible result requires further assumptions before it is generalized to a new population [177].

Let $S_i = 1$ indicate participation of facility $i$ in the study, $A_i \in \{0, 1\}$ a prespecified policy assignment, $Y_i(a)$ a potential outcome in the unit defined by Section 27, and $X_i$ attributes measured prior to selection and assignment. Even if a randomized study identifies $\mathbb{E}[Y_i(1) - Y_i(0) \mid S_i = 1]$, the corresponding target-population mean requires a justified transport procedure. One candidate formula, *only if* trial participation is conditionally independent of the potential outcomes given $X$, the target covariate distribution is measured, the treatment is comparable across settings, and the necessary support exists, is

$$\tau_{\text{target}} = \int \Big\{ \mathbb{E}[Y \mid A = 1, S = 1, X = x] - \mathbb{E}[Y \mid A = 0, S = 1, X = x] \Big\}\, dF_{X|\text{target}}(x). \tag{85}$$

For randomized trial arms, exchangeability and consistency remain necessary; for observational assignments, further confounding assumptions are required. If a target community has production conditions outside the trial's support, Equation (85) relies on extrapolation rather than empirically identified transport. The difference between operational settings cannot be resolved by adding more decimal places to a facility-level efficiency measurement.

Case selection needs a public or access-controlled register of *eligible* cases, not only those finally published. Investigators should count refusals, failed installations, discontinued pilots, and sites excluded for technical reasons; describe baseline capabilities; document target populations; and stratify reporting by geography, production type, system maturity, and governance constraints. Failed attempts to recruit proprietary facilities are evidence about the study's observation coverage, though not proof of misconduct by any refusing actor. A relevant negative result may be that a policy improvement remains valid among highly instrumented cells but cannot be transported to a target population because support is absent. Such an outcome narrows the scope of the research without undoing the local result.

DRAFT

## 29.4 What Incomplete Supplier Networks Allow Researchers to Bound

Production networks are not complete merely because a first-tier supplier list is available. Firms may purchase through intermediaries, substitute vendors following disruptions, use shared infrastructure, and draw on materials processed through several geographic tiers. Multi-tier tracing faces supplier nondisclosure, inconsistent product identifiers, ambiguous functional equivalence, and varying time resolutions. Upstream input–output accounts provide contextual estimates but may smooth over product-specific concentrations and site-level exposures [173, 142, 190]. Link-tracing techniques also require sampling assumptions; evidence from network-sampling methodology cautions against regarding a visible sampled graph as a representative population network by default [48]. The latter is a methodological analogy, not an assertion that respondent-driven sampling estimators directly apply to industrial supply chains.

For an additive burden metric such as kilograms of $CO_2$-equivalent over a declared lifecycle and functional unit, partition the relevant contributions into an observed set $O$ and a *defined, finite* set $U$ of omitted modules. Suppose validated observed contributions sum to $B_O$ and justified, same-unit bounds satisfy $\ell_j \le b_j \le u_j$ for each $j \in U$. Then elementary interval arithmetic gives

$$B_{\text{total}} \in [L, U_B] := \left[ B_O + \sum_{j \in U} \ell_j, \quad B_O + \sum_{j \in U} u_j \right]. \tag{86}$$

This requires consistent and nonoverlapping allocation of the contributing modules, aligned functional units and observation horizons, valid bounds, and an exhaustive accounting of the designated module universe. When the number of relevant omitted tiers is unknown or no defensible upper bounds are available, a finite upper endpoint cannot be asserted. Some nonadditive ecological processes, including threshold responses and interacting exposures, require a process model instead of summing burdens. For a direct burden such as cumulative energy use, a physically justified lower bound can remain useful even when the upper bound is indeterminate.

**Proposition 18** (A decision threshold under bounded omission)**.** *Consider two policies $a \in \{0, 1\}$ supplying the same predeclared service, with total burdens $B_a \in [L_a, U_a]$ within an identical accounting boundary and burden metric. Under no further assumptions about their joint relationship, the contrast $\Delta = B_1 - B_0$ lies in $[L_1 - U_0, U_1 - L_0]$. Consequently, if $U_1 < L_0$, policy 1 has lower burden for every completion compatible with the stated intervals; if the contrast interval contains zero, the sign is not established by the bounds alone.*

*Proof.* The smallest possible difference between two interval-bounded quantities is no smaller than $L_1 - U_0$, and the largest is no greater than $U_1 - L_0$. A strict $U_1 < L_0$ makes the upper endpoint negative. When zero lies inside, one cannot infer a strict sign from the marginal bounds without additional, validated restrictions coupling the policies; some dependencies between completions could narrow the interval. The interval is a conservative outer bound and need not be sharp; even a strict marginal separation warrants an actual policy comparison only if both policy bounds refer to the same functional service, boundary and realizable scenario assumptions. □

The proposition is a mathematical consequence of declared intervals, not evidence that either policy has a lower actual ecological impact. In a hypothetical comparison, suppose one plant reports lower direct electricity consumption under policy 1 while purchasing an additional embodied component whose upstream production is outside the meter boundary. Two supply-network completions can agree on the entire observed meter series yet attribute very different burdens to the unmeasured component. Without a justified bound or further measurement, the observed savings cannot determine the lifecycle sign. Conversely, if independently audited conservative bounds satisfy the proposition's strict separation condition, an investigator may establish a robust burden ordering without claiming to know the missing tiers exactly. This gives partial identification a positive role beyond a generic statement that knowledge is incomplete [118].

DRAFT

Table 52: Partial network observations and defensible inference. Categories distinguish empirical detection from unobserved structural assumptions; thresholds must be declared separately for each case.

| Observed material | Warranted inference | Additional evidence needed for stronger claim |
|---|---|---|
| Metered electricity at one site | Bounded direct electricity use, conditional on meter validity | Electricity sourcing, upstream equipment burdens, system displacement |
| Known first-tier purchase orders | Documented transactions for named components | Supplier genealogy, actual material use, alternatives and buffers |
| Verified maintenance work orders | Recorded interventions and observed downtime | Undocumented labor, common causes, counterfactual failure rates |
| Sampled ecological receptor | Indicator in sampled places and intervals | Pre-intervention trajectory, detection power, spatial spillovers, competing pressures |
| Governance document and interface log | Stated authority and recorded exercise of an interface | Effective veto capacity, enforcement, informal practice and remedy access |

## 29.5 Causal Attribution Across Coupled and Changing Systems

The counterfactual estimands of Section 27 do not become identifiable solely through an expanded dependency graph. The choice of an energy-oriented controller may follow an electricity-price shock that also affects output and carbon intensity. Maintenance interventions can respond to deterioration while improving subsequent availability. A facility's adoption of AI software may coincide with new equipment, workforce retraining, and changed procurement contracts. Contemporaneous changes can therefore generate policy–outcome associations even when their causal roles differ. A target-trial emulation requires an intervention, eligibility rule, allocation strategy, time zero, follow-up, outcomes, and estimand defined before selecting adjustment and estimation methods [54, 160].

Coupled infrastructures add interference. A plant that reduces its grid load may affect marginal generation and prices elsewhere; a major supplier substitution can shift shortages to another purchaser; shared ecological receptors can accumulate exposures from multiple installations. If treatment $A_i$ for unit $i$ influences the potential outcome of unit $j$, a simple no-interference comparison lacks a stable unit-level counterfactual. Under a declared exposure mapping $g_i(\mathbf{A})$, a possible estimand relates $Y_i(a_i, g_i)$ to alternative assignments, but the exposure mapping and its restrictions require empirical justification [56]. An investigator should compare isolated-cell, linked-cluster, and region-level specifications, and inspect whether apparent savings relocate burdens between units. Interference need not invalidate every analysis; it changes what the intervention and estimand mean.

Longitudinal adjustment presents another difficulty. A variable such as component deterioration can affect both the next control policy and subsequent repair cost, while itself being affected by earlier policy settings. Conventional adjustment for such a post-treatment variable may bias estimates of a total policy effect; appropriate sequential methods depend on the causal history and assumptions. Difference-in-differences can be informative where plausible trends and treatment-timing requirements hold, but pretrends and negative-control outcomes are diagnostics rather than self-proving identification. Instrumental-variable approaches require defensible exclusion, relevance, and interpretation assumptions. The design should report why a chosen approach is credible and assess its plausible failures, instead of describing the availability of a technique as proof of causality.

Aggregate evidence requires special restraint. A regional rise in industrial robot installations alongside a change in employment, energy use, or material throughput cannot identify the causal consequences of autonomous AI production without addressing composition changes, macroeconomic shocks, demand, and selection. Historical-materialist interpretation can propose mechanisms relating machinery, accumulation, labor, and institutional control; it does not license the inference that temporal coincidence measures a particular mechanism. An adequate analysis should offer competing explanations, disaggregate outcomes when possible, use appropriate control or comparison histories, and make the distinction

DRAFT

between mechanism interpretation and identified effect explicit.

## 29.6 Ecological Baselines and the Problem of Recoverability

Ecological inference requires more than a lifecycle burden estimate. Ambient temperature, water availability, population structure, habitat function, and recovery trajectories differ across locations and observation periods. A field station established after industrial expansion can document the subsequent condition but may not observe the system that preceded disturbance. Weather, land management, invasive species, other industries, and delayed responses can move the measured indicator independently of the production policy. An inventory result such as kilograms of emissions per functional unit measures a pressure according to its accounting conventions; it is not a direct observation that an ecological community has recovered. The ecological framework of Sections 21–24 calls for specifying receptors, thresholds, monitoring, spatial connections, and restoration hypotheses separately from input efficiency [84, 83, 59, 129].

A defensible comparison begins with a causal account of the ecological indicator, its temporal and spatial support, measurement error, and plausible exposure lags. Repeated pre-intervention sampling, reference sites, independently validated remote observations, and process-based sensitivity analyses can improve evidence. These procedures remain imperfect where the unimpacted reference state is unobservable, historic landscapes have already changed, or climate trends shift the conditions of comparison. The absence of detected deterioration is informative only together with detection limits and the minimally important change that the study could have found. Likewise, a monitored increase in a selected biodiversity indicator need not demonstrate recovery of all ecological functions or affected nonhuman interests.

Robust viability, as formalized in Section 24, depends on declared constraints, uncertainty sets, controllable actions, and dynamics. A feasible policy in a simplified model is therefore a conditional existence result. If ecological process parameters cannot be identified from the monitoring record, its viability classification can be sensitive to the choice of dynamics or uncertainty radius. Investigators should report which constraints are binding, whether feasible sets persist across defensible process models, and which missing data could reverse the conclusion. Where species loss, aquifer change, or contamination cannot be reliably offset, the decision problem may involve a non-substitutable limit; it cannot be settled by assigning an unexamined exchange rate between environmental loss and economic output.

## 29.7 Competing Dependency Models and the Test of Additional Explanatory Value

The proposed generative-relational (GR) representation joins typed processes, participants, enabling inputs, authority, provenance, and ecological receptors. Its explanatory capacity depends on how these categories are defined, connected, and measured. Different valid representations can encode the same evidence. A component-level bill of materials, an input–output table, a temporal production network, a causal directed acyclic graph, and an agent–activity–entity provenance graph address overlapping but nonidentical tasks. None acquires causal authority from its graphical form. In particular, an edge recording purchase history should not be interpreted as a physical necessity; a path across suppliers should not be interpreted as a causal effect; and a model's ability to fit an observed history does not establish predictive stability after reorganization [112, 196, 152].

Model uncertainty should be studied at several levels: (i) ambiguity in defining an activity and a boundary, (ii) uncertainty as to whether a relation exists, (iii) misclassification of its type or direction, (iv) uncertain parameters of dependency or ecological response, and (v) incompatible structural explanations for the same observed events. Investigators can vary node granularity, substitution structure, environmental boundary, and intervention assumptions, recording which practical conclusions remain invariant. Where an input may be indispensable under one network completion and redundant under another, the appropriate conclusion is model dependence until new intervention or supplier evidence distinguishes them. The earlier typed formalism offers a notation for the problem, not a guarantee of correct structure.

DRAFT

The benchmark proposed in Sections 27 and 28 remains essential. GR-augmented analyses must receive the same underlying evidence as strong LCA/MFA, time-indexed production-network, reliability, and provenance baselines, with appropriate permission to include authority and event-history attributes. The incremental question is whether GR's particular relational typing or inference workflow improves a prespecified task after the baseline has access to equivalent substantive information. Blind evaluations should use held-out facilities or time-separated disruptions, adjudicated event labels, predictive calibration, analytical workload, storage and computing costs, and sensitivity to withheld tiers. A favorable qualitative case with a deliberately weak comparator would not suffice. A repeated null advantage, poorer calibration, or avoidable complexity is a legitimate negative result; it limits the method's incremental utility without proving that material relations are irrelevant.

## 29.8 Power-Conditioned Observability and the Interpretation of Silence

The relation between power and evidence is a necessary empirical question, not a universal explanation for missing records. Owners, platform vendors, public agencies, trade associations, workers, subcontractors, and affected communities may possess very different abilities to create, preserve, request, challenge, and interpret production records. Commercial secrecy and data-security constraints can be legitimate, yet they may also produce systematic gaps around subcontracted labor, proprietary maintenance interfaces, upstream material processing, or localized ecological harms. Equally, organizationally produced records may overrepresent easily counted throughput and operational uptime while underrepresenting care work, informal repair, and displaced exposure. These possibilities follow from the institutional analysis in Section 26 and warrant direct investigation rather than presumed motives [153, 146, 74].

An inquiry can operationalize observability as a vector of auditable rights and practices: who is entitled to record an event; who has access to raw measurements; whether an affected party can contest a classification; who may authorize independent verification; what is retained after a contract ends; and what remedies exist for material misreporting. These should be compared with actual outcomes of access requests and sampling, not inferred from an ownership label alone. If facilities with different governance arrangements yield systematically different access after accounting for technology and applicable rules, that finding can motivate a narrower causal investigation. Conversely, comparable disclosure practices across organizational forms would count against a simple power-determines-visibility explanation. This approach treats power as a potentially explanatory condition rather than a substitute for observation.

Independent auditing requires safeguards. A public repository of precise industrial-network data may expose commercially sensitive information, worker identities, critical infrastructure, or ecological locations needing protection. Controlled researcher access, provenance-preserving redaction, tier-aggregated summaries, confidential third-party verification, and publication of measurement instruments can sometimes support scrutiny without indiscriminate disclosure. None of these methods guarantees full transparency or resolves who is entitled to the underlying data. Conflicts between privacy, security, scientific verification, and the standing of affected communities call for institutionally legitimate decisions. The important analytical distinction is between a relation shown not to exist and a relation that the observer was *not empowered or equipped to see*.

## 29.9 Registering Disconfirmation Across the Research Programme

The paper's main propositions should be distinguished from their most convenient supporting anecdotes. A meaningful test specifies an empirical target, a competent comparator, a predicted contrast or pattern, and the result that would restrict the claim. Table 53 proposes a prospective register. Its entries are criteria for future studies rather than tests already performed. For causal hypotheses, failure of identification is marked separately from evidence of a null or reversed effect; for the formal model, a mathematical contradiction under the same assumptions is distinct from an empirical discrepancy with its application.

These tests should not be phrased so that every possible result is interpreted as confirming a theory. An observed absence of concentrated suppliers is evidence against concentration in the sampled setting; it does not automatically turn into confirmation of a subtler hidden concentration. A demonstrated

DRAFT

Table 53: Proposed prospective register for principal hypotheses. Each outcome concerns a declared boundary, population, and horizon; missing identification is distinct from disconfirmation.

| Testable proposition | Competent test or comparator | Finding that limits the proposition |
|---|---|---|
| Operational automation leaves material reproduction dependencies | Audited bills of materials, service continuity, substitute and maintenance inventories | Dependencies materially decline under stable expanded boundaries; proposed links are not functionally necessary |
| Automation can relocate rather than remove bottlenecks | Matched pre/post network and interruption analysis; independent vendor records | Diverse substitutes expand effective continuity and no offsetting critical concentration is supported |
| Institutional rights condition executable choices | Verified feasible-set and decision audits across comparable technical settings | Differences in rights or access fail to explain policy capability or choices after competing constraints are tested |
| Reorganization alters labour and social-reproductive burdens | Worker-time, skills, exposure, quality, and household/service indicators with proper comparison groups | Proposed displacement or invisibility fails to appear, or benefits predominate within declared measures |
| Unit efficiency does not establish absolute ecological relief | Extended LCA and system-response estimates including displaced demand and suppliers | Robust absolute reductions survive credible upstream and scale-response bounds |
| A GR representation adds analytic value | Strong temporal-network, LCA/MFA and reliability baselines with equal information access | No repeatable benefit in calibration, diagnosis or decision quality relative to cost |

environmental improvement is evidence against a claim that the particular intervention merely exports burdens, even if other systems might exhibit displacement. A strong conventional model may reproduce every relevant trajectory feature with fewer assumptions than GR. The response to such findings should be to narrow the hypothesis, improve the observation design, or abandon an unproductive component. Retrospectively selecting a new threshold or observation horizon to preserve a favored conclusion would defeat the registration.

Null results also require adequate measurement power. A wide uncertainty interval compatible with substantial benefit and harm is a statement of unresolved evidence, rather than strong confirmation of no effect. Selective publication of favorable facilities, outcome switching, and repeated unregistered model searches can undermine the evidentiary value of apparent successes; transparent protocols and full study registers are therefore relevant beyond conventional biomedical settings [102]. This citation concerns a methodological analysis of bias and selective publication; it supplies no empirical estimate of the proportion of false findings in autonomous-production research. If a data steward cannot permit publication of the necessary negative results, the resulting case may be unsuitable for a confirmatory claim. The empirical programme should publish its coverage limitations and rejected hypotheses alongside any positive outcomes.

## 29.10 A Reporting Protocol for Research and Public Deliberation

An integrated report should carry a sequence of named decision points. First, freeze the version of the productive service, time horizon, site eligibility, environmental boundary, and relevant rights and constraints. Second, disclose the measurement instruments and distinguish observed, imputed, modeled, withheld, and unknown records. Third, register the analytic estimand and its causal, sampling, physical, and functional assumptions. Fourth, compare alternative completions of supplier, maintenance, ecological, and authority networks. Fifth, quantify the robustness of each consequential conclusion, using partial-identification intervals where point identification is unavailable. Sixth, test whether another competent method offers a simpler explanation with similar or superior performance. Seventh, expose the distributional and normative decisions not answered by the empirical model.

For public deliberation the report should separate three outputs. The *evidentiary output* states what was observed, modeled, or bounded. The *causal output* states which differences between policies or institutions can be defended given their identification assumptions. The *normative output* identifies whose ecological and material interests are at stake, which limits have legal or ethical standing, what remains contested, and who has authority to decide. Such separation does not sever the connection between evidence and values: a measured harm matters precisely because of the lives, livelihoods, and ecological processes implicated. It instead prevents uncertainty from being concealed inside an optimizer or val-

DRAFT

Table 54: Minimum disclosure before advancing a comparative-production claim. Each element concerns the specified study rather than a universal completeness requirement.

| Gate | Required report | Reason to withhold a stronger conclusion |
|---|---|---|
| Observation | Site, service, instrument, access, unit, missingness and provenance | Unreported denominator or non-ignorable nonresponse |
| Identification | Estimand, treatment, comparison, assumptions and interference model | Confounding, missing overlap, undefined counterfactual |
| Boundary | Supplier and ecological coverage; modeled and unbounded tiers | Plausible omitted flows change the claimed direction |
| Robustness | Alternatives, perturbations, prediction validation, partial bounds | Sign or practical choice reverses under credible scenarios |
| Normative status | Affected parties, thresholds, authority, contestation and remedies | Model rankings presented as authorized ethical outcomes |

ues from being misrepresented as scientific coefficients whose legitimacy follows from computational feasibility.

The research programme also benefits from staged commitments. A first-stage audit of one industrial site may legitimately end at service logs and material balances; a second stage may validate substitution and interruption models; a third may extend to representative supplier and receptor sampling; only later can a cross-institutional or regional causal comparison become credible. Each stage should have stopping rules when unknown inputs, poor overlap, high measurement error, or incompatible ecological models make the next inference unwarranted. Independent review should draw on manufacturing engineering, industrial ecology, causal inference, labour research, ecological field science, and institutional analysis. Such review is a methodological requirement for stronger future claims, not an assertion that the current manuscript has undergone it.

## 29.11 What the Inferential Limits Establish

The section establishes a disciplined boundary between a represented trajectory and an inference about productive reality. Evidence can be incomplete because events were never recorded, because collection and access depend on institutions, because chosen sites differ from the intended population, or because competing causal and ecological models fit the same observations. These conditions have distinct remedies: improved measurement, access governance, better sampling, stronger intervention designs, sensitivity analysis, bounds, or a candid statement of non-identification. The interval proposition shows a narrow but useful positive result: some comparisons remain decisively ordered under defensible bounds, while others cannot support a directional conclusion. No general ecological or institutional verdict follows from the formalism alone.

The candidate GR framework remains a provisional way to coordinate these inquiries. Its value is conditional on accurate typed records, defended assumptions, and better performance for declared tasks than equally informed conventional methods. Its limitations include observation bias, ambiguously typed dependencies, uncertainty about causal structure, and the inability of relational bookkeeping to authenticate events or decide normative standing. These are scientifically consequential rather than merely editorial restrictions. In some settings the research may find that conventional industrial ecology and supply-network models suffice, or that automation reduces critical dependence and absolute ecological pressures more effectively than originally hypothesized. Such results must remain admissible.

The concluding section can therefore synthesize the paper on a narrower and more defensible basis: increasingly autonomous AI-mediated production may change the distribution of labour, material dependencies, institutional authority, and ecological burdens, but their direction and viability are empirical, historically situated questions. The necessary continuation is a staged, independently reviewable research programme that links productive service, physical flows, social reproduction, ecological conditions, and contestable decisions without collapsing technical capacity into self-reproduction, or empirical

DRAFT

description into normative entitlement.

# 30 Conclusion and Research Continuations

This section brings the historical, formal, engineering, ecological, and comparative inquiries into a single account of the conditions under which increasingly autonomous AI-mediated production might sustain its own productive capacity while remaining compatible with the continuation of affected ecological processes. The investigation began with an asymmetry between what a production system can accomplish during an observed operating interval and what is required to sustain the material and social conditions of further operation. Its conclusion is deliberately conditional. Greater automation changes productive possibilities and can reorganize dependencies, but the direction of those changes, their distributive effects, and their ecological consequences cannot be deduced from an automation label. They require measurements over justified system boundaries, explicit institutional comparisons, and observations appropriate to the ecological processes at issue. The present manuscript supplies a framework and a research design for these inquiries, alongside sourced evidence from existing literature. It does not claim a demonstrated self-reproducing AI production system, an original estimate of ecological benefit, or a generally superior form of ownership.

## 30.1 The Historical Problem of Reproducing Productive Capacity

The paper's historical-materialist starting point is that productive capacity is constituted through relationships among workers, knowledge, machinery, material supplies, infrastructures, organizational practices, and the social conditions through which these capacities are reproduced. The historical incorporation of collective knowledge into machinery, the development of industrial automation, and the contemporary diffusion of AI-supported planning and control do not remove these relationships. They can transform which participants perform particular tasks, where necessary inputs are procured, which actors possess effective decision rights, and what investments become prerequisites for production. This analytical orientation follows the tradition of examining machinery and the reproduction of social production, without equating a present technical architecture with a prediction about the next historical order [120, 125, 124].

The historical analysis and typed dependency model establish why technical performance, renewal of productive inputs, control over access, and social reproduction require separate records. Across the manuscript, *operational autonomy* refers to bounded task execution, *reproductive autonomy* to the renewal of enabling capacities under recorded support and horizon conditions, *ecological compatibility* to specified and justified ecological endpoints, *institutional implementability* to effective access and authorization of a feasible policy, and *legitimacy* to a further contestable judgment about the authority and procedures governing that policy. They are linked propositions, not interchangeable labels or entries on a universal autonomy scale. A workcell can operate with few on-site interventions while relying on externally produced parts, specialist labour, and restricted knowledge for the next production interval. These dependencies are mediated through wages, procurement, contracts, public infrastructure, organizational control, and material exchange. Their reorganization is therefore a question about productive forces *and* the relations through which those forces can be maintained and used, rather than a numerical measure of task automation alone.

This reasoning also limits the paper's interpretation of scarcity and productive power. A dependency becomes a consequential bottleneck relative to a specified service, alternative pathways, and horizon. Institutional power requires the further possibility of exercising control over access or adjustment. Neither high technical centrality nor ownership alone guarantees that control will be effective. Conversely, an apparently dispersed production network can remain dependent on shared infrastructures or permissions. The proposed investigation of *dependency relocation* therefore concerns empirical changes in the composition and substitutability of requirements, not a presumption that automation inevitably produces concentration or extracts value. Labour reorganization must similarly be investigated through documented

DRAFT

changes in tasks, skill, remuneration, care, maintenance, and social reproduction rather than inferred from the adoption of a machine [43].

These distinctions support historically situated comparisons. They leave open the possibility that automation can diversify a supply network, improve worker safety, lower burdens, or expand socially useful productive options. They also leave open the possibility of geographically displaced costs, tightened access to critical equipment, intensified extraction, or increased demands on unrecognized labour. Each mechanism admits different observations and different potential refutations. The resulting account of productive transformation is comparative rather than historically predetermined.

## 30.2 What the Proposed Relational Representation Contributes

The generative-relational (GR) representation introduced in Section 6 is intended to make heterogeneous productive requirements and their changing relations inspectable across boundaries. Its records distinguish entities and events, material and energy quantities, service capacities, access conditions, evidence provenance, and ecological attributes. Their types matter. A measured electrical input in kilowatt-hours, a functional requirement for an equipment component, a legal permission, a report of possible contamination, and a claim about legitimate authority cannot be combined into one scalar quantity without an additional and justified model. The representation offers a common interface for comparing the relationships among these observations, not a replacement for their domain-specific measurement disciplines.

The paper's formal contributions are conditional and have different empirical implications. The distinction between operational and reproductive autonomy in Section 7 requires researchers to test continued capacity, not only completion of a campaign. The upstream-requirements constructions in Section 8 require appropriate flow coefficients, substitution assumptions, and evidence on missing supplier tiers. The policy-set inclusion in Section 26 requires observation of financing, permissions, and enforceable decision rights before it can inform a claim about realized production. The robustness calculations in Section 24 are conditional on the stated ecological-state model, constraints, disturbance set, and treatment of uncertain fixed parameters. Each result specifies what would have to be observed for application; mathematical consistency alone establishes none of those observations.

The methodological value of these distinctions depends on comparison with competent alternatives. Material-flow analysis, input–output methods, lifecycle assessment, reliability engineering, supply-network models, established provenance standards, and conventional causal inference already provide substantial theoretical and empirical instruments. The GR framework warrants retention only where a typed, historically changing record supplies additional information for a declared task at acceptable cost and error. Such tasks may include following the interaction between equipment dependencies, replacement histories, ecological exposure, and authorization decisions, provided these connections can be measured and independently checked. An equal-information comparison in which established methods perform equally well or better is an admissible negative result. The framework should then be simplified, restricted to a documentation role, or replaced for that task. Relational vocabulary by itself provides neither an authenticated event history nor privileged epistemic authority.

## 30.3 The Conditions Connecting Technical and Ecological Reproduction

The engineering review in Sections 13–20 differentiates the demonstrated operation of specific robotic, laboratory, and control components from the much stronger capacity to renew their enabling machinery and inputs. Computing facilities, energy and thermal infrastructure, specialized equipment manufacture, spare parts, calibration, repair, and materials recovery occupy different technological maturity levels and follow different replacement cycles. An operating installation may automate one sequence while relying on extensive human and infrastructural support elsewhere. Proposals concerning robotic or AI-supported self-maintenance are consequently bounded engineering hypotheses until the relevant systems have been tested, including disturbances, repair, material replenishment, and longer operating horizons.

Sections 21–25 develop a corresponding distinction on the ecological side. A production regime's ability to continue acquiring its inputs and operating its equipment does not establish that receiving

DRAFT

ecosystems can continue their characteristic processes. Physical inventories describe withdrawals, emissions, stocks, and flows within declared boundaries. Impact assessments connect some of those pressures to selected ecological consequences under further models. Recovery studies concern the response of specific ecological processes after disturbance, which need not be captured by an annual emissions inventory. Robust viability models then ask a narrower formal question: whether admissible production policies satisfy specified service and ecological constraints over a stated horizon and set of disturbances. A mathematically nonempty viability set is conditional on these specifications. It does not certify the sufficiency or legitimacy of the constraints chosen. Likewise, policy implementability within a documented institution is a separate issue from both ecological compliance and the justification of that institution's authority.

The conjunction of productive and ecological continuation is therefore a research problem with at least four separately justified components: a measurable productive service; a defended inventory of its direct and indirect requirements; adequate evidence about affected ecological receptors and their dynamics; and publicly or otherwise legitimately authorized constraints on permitted interventions. Compatibility requires that the selected policy satisfy the relevant constraints under the stated uncertainty model. Where the evidence is incomplete, conclusions may be bounded or indeterminate. Where no admissible policy satisfies the declared requirements, the appropriate result is a documented incompatibility *within that model*, prompting scrutiny of demand, technology, distribution, timing, or the justified requirements themselves. This does not entail that ecological interests should be weakened to recover mathematical feasibility.

Cumulative and induced-scale effects matter because lower unit burdens do not ensure lower total burdens. A more efficient process can coincide with increased total activity, relocated production, or earlier equipment replacement, while other interventions may reduce defects and absolute material use. Section 25 treats these mechanisms as competing possibilities. A reliable comparison must report both service-normalized and absolute outcomes, specify its counterfactual, and examine the time and location of consequences. A locally lower demand on electricity, water, or materials cannot alone support a conclusion about regional ecosystems, distant supply chains, or ecological recovery.

## 30.4 The Evidentiary Status of the Paper's Principal Claims

A clear account of what this manuscript establishes is necessary because its conceptual, mathematical, and empirical components have different kinds of warrant. Table 55 summarizes these distinctions. The categories are not stages on a universal ladder: documented engineering operation may be well established for a narrow task while the causal ecological effect of its introduction remains unidentified; a mathematically sound proposition may be uninformative if its assumptions do not hold in the target setting.

The reported external evidence is likewise source-bounded. Existing reports document selected industrial installations, laboratory demonstrations, energy and materials trends, and occupational exposure indicators. Those observations motivate research questions; they are not a pooled causal experiment on autonomous production. Sections 27–29 set out the further requirements of credible controls, sufficient network coverage, calibrated ecological models, qualified inference, and transparent uncertainty. In particular, observing that a facility has a sophisticated control architecture establishes neither complete supply-chain reproduction nor an identified contribution to ecological recovery.

The distinction between a negative result and an unidentified result is especially consequential. An adequately measured intervention that fails to reduce a prespecified total ecological burden contradicts a directional improvement hypothesis for its tested scope. An intervention with unobserved upstream inputs and no defensible bounds instead leaves the total contrast unresolved. Similarly, a finding that critical suppliers become more substitutable challenges a hypothesis of tightening dependence in that setting, whereas undisclosed supplier identities alone cannot demonstrate hidden concentration. These rules protect the programme against interpreting every outcome as confirmation. The interval-comparison result of Section 29 offers a limited positive possibility: some burden orderings can be defended when con-

DRAFT

Table 55: Claim status at the completion of this exploratory manuscript. “Established” refers only to the stated conceptual or conditional mathematical result, not to empirical verification of a proposed production regime.

| Research object | Warrant available in this paper | What remains unestablished |
|---|---|---|
| Operational and reproductive autonomy | Explicit distinct definitions and conditional examples | Their gap in a specified population of facilities |
| Changed productive dependencies | Typed comparison and bounded accounting procedures | Direction and magnitude under a real intervention |
| Institutional production choices | Conditional separation of physical and implementable policy sets | Which governance arrangement produces better outcomes |
| Ecological continuation and viability | Existing measurement methods and model-dependent feasible-set analysis | Regionally validated recovery and total impact effects |
| Comparative policies and GR utility | Prospective causal designs, evidence rules, and falsification tests | Original causal estimates or incremental predictive benefit |

Table 56: Proposed order of empirical work. Every gate applies to a specified population, productive service, system boundary, observation interval, and set of potentially affected participants or ecological receptors. No stage is reported as completed here.

| Stage | Primary question | Minimum study product | Rule for progression |
|---|---|---|---|
| I. Record | What is demonstrably produced and consumed? | Provenance-checked service logs, input and output accounts, missing-data register | Reconcile units, inventories, and key source records |
| II. Test | Which inputs are functionally consequential? | Interruption, substitution, reliability, and repair tests against competent engineering models | Quantify alternatives and uncertainty; avoid causal claims from topology alone |
| III. Trace | Where do burdens and capacities travel? | Extended supplier flows, labour records, lifecycle inventory and receptor-based ecological observations | Document or bound consequential omissions and transport assumptions |
| IV. Compare | What changes under a defined intervention? | Matched-service prospective or defensibly identified observational comparison with negative-result reporting | Establish overlap, stable intervention definitions, sensitivity, and appropriate interference analysis |
| V. Deliberate | Which feasible trajectories merit authorization? | Transparent options, material consequences, accountable decision records and review opportunities | Distinguish scientific constraints from legitimate institutional choice |

servative ranges remain separated, but overlapping ranges generally do not identify a direction without more information.

## 30.5 A Staged Programme for Independent Empirical Evaluation

The next research step should consist of a staged programme whose successive claims depend on recorded evidence rather than on the ambition of the overall theory. Each stage has its own feasible deliverable and a rule for declining stronger inference. The measurement modules introduced in Section 28 supply possible settings: a qualified autonomous laboratory, a conventional or highly automated production cell, a computing-intensive production supply chain, a maintenance network, and a material–ecological supply network. Selection should be driven by data rights, technically comparable service definitions, and the prospect of valid measurement rather than by a preference for unusually successful demonstrations.

At Stage I, investigators would register a functional unit and its quality requirements, a site and supplier boundary, and a finite observation horizon; audit meters, logs, and inventory records; and identify unresolved links. An honest finding that input records cannot yet be reconciled is a useful research re-

DRAFT

sult. Stage II would identify which dependencies are causal or functional by using qualified interruptions, component replacement, external validation of reliability models, or natural changes with defensible assumptions. Interventions must respect workplace safety and ecological protection; experimental access cannot be treated as a license to impose hazardous disruptions. Stage III would connect the site to physical upstream flows, affected workers and institutions, and relevant ecological receptors. Its purpose is to test whether claimed benefits survive the accounting boundary and whether exclusions can plausibly reverse the conclusion.

Stage IV would evaluate policies under a matched productive service and decision horizon. As developed in Section 27, throughput-, energy-, and ecology-oriented controls need a well-operated conventional comparator, prespecified outcomes, and an assignment or identification strategy appropriate to interference and changing market demand. The protocol should jointly report site-level and wider-system results, with uncertainty and sensitivity to unrecorded flows. Comparative institutional evaluation can be introduced where effective policies, financing, workforce conditions, public constraints, and ecological baselines are sufficiently comparable. A weak comparison between one exceptionally resourced public facility and a different low-investment private plant would otherwise conflate institutional governance with equipment quality and selection.

At Stage V, measured alternatives would be made available for legitimate judgment. Scientific research can constrain an options set by identifying likely harms, feasible actions, unknown consequences, and conflicting requirements. It does not determine the authority of any actor to trade workers' livelihoods against distant habitat damage or to fix ecological thresholds without deliberation. This stage therefore requires explicit affected-party definitions, institutional procedures, lawful data access, accessible explanations, and opportunities for revision where appropriate. The relationships between cross-border evidence, trust, and remedial authority exceed this paper's scope and remain reserved for a companion inquiry.

This staged sequence need not be implemented in a single giant study. Independent research teams may advance particular modules at different speeds and with differing geographic coverage. Their evidence should remain interoperable through reported units, boundaries, identifiers, provenance, and uncertainty rather than forced into a universal integrated score. A consortium might legitimately complete high-quality facility-level accounting without being able to estimate regional ecological outcomes. Research reports should state that limitation without treating it as a failure to produce useful knowledge.

## 30.6 Comparative Institutions and the Governance of Uncertainty

The material possibilities of autonomous production are shaped by decisions about investment, data, equipment access, labour organization, technology licensing, standards, and the ecological conditions under which production may proceed. Section 26 therefore compares institutionally available production trajectories without positing that capitalist, public, cooperative, or commons-based arrangements possess an automatic ecological or emancipatory property. Each can create both capabilities and constraints. Effective accountability depends on specific decision rules, material capacity, enforceable duties, and the possibility of evaluating consequences.

An institutional assessment should keep three issues separate. *Capacity* concerns what an organization can actually implement, finance, maintain, and monitor. *Distribution* concerns how productive benefits, access, burdens, and future capabilities are allocated among participants and affected others. *Legitimacy* concerns whether decisions affecting those participants and ecological interests are authorized and open to appropriate challenge and revision. Evidence of high output or low direct cost may illuminate capacity while revealing little about distribution or legitimacy. Similarly, an institutionally inclusive procedure cannot substitute for valid measurements of physical harm. Responsible comparison must be able to find different strengths, contradictions, and failures within the same arrangement.

Institutional control can also affect the visibility of evidence. Supplier disclosures, worker reports, repair logs, ecological monitoring, and publication rights are produced under historically specific governance conditions. This observation motivates investigations of selection and missingness, not an un-

DRAFT

supported claim that a particular actor deliberately suppresses evidence. Where strategically withheld information is hypothesized, the research requires observations capable of distinguishing it from confidentiality obligations, fragmented recording systems, or technical absence of measurements. A sound institutional comparison should make the conditions of observation part of its design, including data-access constraints that can prevent a confirmatory claim.

The ethical perspective remains plural. Ecological viability addresses constraints on certain physical and biological processes; social reproduction addresses conditions for human lives and labour; legal and political evaluation concerns authority, rights, voice, and remedy. Relations among these interests can be contested, and no single numerical objective has inherent authority to resolve them. The contribution of the present research is to make the consequences and assumptions of proposed choices more inspectable so that disagreement need not rest on hidden material accounts or mislabeled scientific certainties.

## 30.7 Conditions That Would Restrict the Research Programme

A continuing theory benefits from specifying what would reduce its scope or warrant, not only what might extend it. The proposed research programme should be restricted when one or more of the following conditions holds within an adequately measured, declared setting: automation decreases external input intensity and system-wide absolute burdens without consequential displacement; supply substitution becomes easier rather than more concentrated after an intervention; improved controls achieve ecological and productive requirements together without significant tradeoffs under tested disturbances; or the claimed GR representation adds no material diagnostic or predictive value beyond established comparators. These findings are compatible with portions of the paper's general problem statement while contradicting particular directional hypotheses. Two elementary constructed cases illustrate the required discipline. First, if impact intensity falls from $10$ to $8$ kg per qualified unit while output rises from $100$ to $150$ units, total registered impact rises from $1{,}000$ to $1{,}200$ kg despite the improvement per unit; the larger total could alternatively be reduced by a different demand response. Second, a regime with a larger implementable policy set may choose a higher-burden option than a regime with fewer choices. Neither arithmetic example predicts an actual factory or ranks institutional forms. Together they require that intensity, total scale, available options, selected action, and measured ecological responses be tested as separate claims. They should not be recast as proof of concealed adverse effects merely because such effects remain imaginable.

Other outcomes imply suspension rather than rejection of an empirical claim. If supplier coverage is too low to bound impacts, if policy groups have no credible overlap, if ecological response models have been extrapolated outside their domain, or if proprietary data cannot be independently audited, the corresponding effect may be unidentified. At that point the appropriate product is a coverage report, sensitivity range, or study redesign. The absence of a defensible causal estimate is not evidence that the effect equals zero. These distinctions follow the identification and sensitivity requirements specified in Section 29.

A further limitation lies in the representations themselves. The manuscript treats decision intervals and assessment horizons as modeling devices, not as an unexplained universal ontology of historical time. Changes in ecological conditions may have long memory; work, infrastructure, and supply chains may follow asynchronous cycles; institutional rules can change during observation. If a chosen state representation omits these histories, its transition or viability calculations may be misleading even when internally consistent. Model revision may therefore require new state variables, non-Markovian dynamics, historical interpretation, or different research questions. The formalism is useful insofar as it discloses these assumptions and remains answerable to observation.

Before any claim of established scientific novelty or submission-ready reliability, a full independent review is required. The review should verify original source metadata and interpretations, formal proofs and units, accounting boundaries, engineering maturity descriptions, ecological response assumptions, and institutional and labour claims. The current bibliographic and LaTeX checks establish compilation and internal reference integrity, not scholarly verification of each source or assertion. The manuscript's

DRAFT

proposed protocols are reproducibility aids rather than records of completed experimental implementations.

## 30.8 The Scope of the Resulting Research Contribution

The contribution is an integrated, testable account of a historically situated problem: technological independence within a bounded operating task may coexist with greater dependence on the wider reproduction of productive and ecological conditions. The paper makes five research transitions explicit: from task output to the restoration of productive capacity; from local inputs to qualified upstream alternatives; from observed dependencies to exercisable institutional power; from measured physical flows to causal ecological consequences; and from technically feasible policies to legitimately selected ones. Each transition calls for additional records, assumptions, and comparators, rather than following automatically from the preceding claim. The proposed GR representation is one candidate means of keeping the transitions auditable, subject to equal-information evaluation against established methods. Its strongest prospective claim is therefore conditional and discriminating: a reproducible improvement on defined cross-domain tasks, relative to equally informed and adequately engineered alternatives, with uncertainty and investigation costs declared. If that test fails, the conceptual distinctions may remain useful while a separate claim of methodological superiority should be withdrawn.

The paper thus establishes conceptual distinctions, conditional models, and research designs rather than a demonstrated trajectory toward fully self-reproducing AI production. It makes no deterministic prediction about employment, institutional order, or ecological benefit. A documented improvement in productive independence, ecological continuation, or institutional accessibility would count as evidence for the corresponding bounded claim; documented deterioration would count against it. Where alternatives are observationally indistinguishable, the finding is unresolved rather than favorable by default. On these terms, ecological compatibility is a researchable, historically organized possibility whose attainment would require independent technical, ecological, and institutional evidence.

Future research should begin with independently auditable service, maintenance, supplier, and ecological records; proceed to matched interventions and sensitivity analysis; and finally support decisions through accountable institutions with explicit rights of correction and remedy. The research architecture is justified only where this sequence produces better warranted conclusions or more usable, contestable evidence than competent alternatives. Its remaining scientific and ethical questions are accordingly left open to external tests and deliberation.

# References


[1] Omar Aboul-Enein et al. *Continuous Mobile Manipulator Performance Measurement Data*. 100-57. National Institute of Standards and Technology, 2024. DOI: 10.6028/NIST.AMS.100-57.

[2] Daron Acemoglu and Pascual Restrepo. “Automation and New Tasks: How Technology Displaces and Reinstates Labor”. In: *Journal of Economic Perspectives* 33.2 (2019), pp. 3–30. DOI: 10.1257/jep.33.2.3.

[3] Daron Acemoglu and Pascual Restrepo. “Robots and Jobs: Evidence from US Labor Markets”. In: *Journal of Political Economy* 128.6 (2020), pp. 2188–2244. DOI: 10.1086/705716.

[4] Philippe Aghion and Jean Tirole. “Formal and Real Authority in Organizations”. In: *Journal of Political Economy* 105.1 (1997), pp. 1–29. DOI: 10.1086/262063.

[5] Aaron D. Ames et al. “Control Barrier Function Based Quadratic Programs for Safety Critical Systems”. In: *IEEE Transactions on Automatic Control* 62.8 (2017), pp. 3861–3876. DOI: 10.1109/TAC.2016.2638961.

[6] ASML Holding N.V. *Annual Report 2025*. 2026. URL: https://investor.asml.com/annual-reports/2025 (visited on 10/26/2026).

[7] Jean-Pierre Aubin. *Viability Theory*. Boston: Birkhäuser, 1991. URL: https://viability-theory.org/en/node/51.

[8] Albin Bajrami et al. “Investigating Collaborative Robotic Assembly: A Case Study of the FANUC CRX-10 iA/L in Industrial Automation at i-Labs”. In: *Eng* 5.2 (2024), pp. 532–543. DOI: 10.3390/eng5020029.

[9] Richard Baldwin and Rebecca Freeman. “Risks and Global Supply Chains: What We Know and What We Need to Know”. In: *Annual Review of Economics* 14 (2022), pp. 153–180. DOI: 10.1146/annurev-economics-051420-113737.

[10] Ken Ball. *The Dawn of the Programmable Logic Controller (PLC)*. Aug. 2, 2015. URL: https://www.automation.com/article/the-dawn-of-the-programmable-logic-controller-plc (visited on 10/26/2026).

[11] Christophe Béné, Luc Doyen, and David Gabay. “A Viability Analysis for a Bio-Economic Model”. In: *Ecological Economics* 36.3 (2001), pp. 385–396. DOI: 10.1016/S0921-8009(00)00261-5.


DRAFT

[12] Julia Bitencourt, Ana Wooley, and Gregory Harris. “Verification and validation of digital twins: a systematic literature review for manufacturing applications”. In: *International Journal of Production Research* 63.1 (2025), pp. 342–370. DOI: 10.1080/00207543.2024.2357741.

[13] Franco Blanchini. “Set Invariance in Control”. In: *Automatica* 35.11 (1999), pp. 1747–1767. DOI: 10.1016/S0005-1098(99)00113-2.

[14] Daniil A. Boiko et al. “Autonomous Chemical Research with Large Language Models”. In: *Nature* 624 (2023), pp. 570–578. DOI: 10.1038/s41586-023-06792-0.

[15] Robert Brenner. “Agrarian Class Structure and Economic Development in Pre-Industrial Europe”. In: *Past & Present* 70.1 (1976), pp. 30–75. DOI: 10.1093/past/70.1.30.

[16] Benjamin Burger et al. “A Mobile Robotic Chemist”. In: *Nature* 583 (2020), pp. 237–241. DOI: 10.1038/s41586-020-2442-2.

[17] Brantly Callaway and Pedro H. C. Sant’Anna. “Difference-in-Differences with Multiple Time Periods”. In: *Journal of Econometrics* 225.2 (2021), pp. 200–230. DOI: 10.1016/j.jeconom.2020.12.001.

[18] Martina Calzavara et al. “Achieving Productivity and Operator Well-Being: A Dynamic Task Allocation Strategy for Collaborative Assembly Systems in Industry 5.0”. In: *The International Journal of Advanced Manufacturing Technology* 134 (2024), pp. 3201–3216. DOI: 10.1007/s00170-024-14302-3.

[19] Richard Candell et al. *Guide to Industrial Wireless Systems Deployments*. NIST AMS 300-4. National Institute of Standards and Technology, 2018. DOI: 10.6028/NIST.AMS.300-4.

[20] Richard B. Canty and Milad Abolhasani. “The Past, Present and Future of Self-Driving Laboratories”. In: *Nature Reviews Chemistry* 10 (2026), pp. 523–537. DOI: 10.1038/s41570-026-00847-2.

[21] Vasco M. Carvalho. “From Micro to Macro via Production Networks”. In: *Journal of Economic Perspectives* 28.4 (2014), pp. 23–48. DOI: 10.1257/jep.28.4.23.

[22] Vasco M. Carvalho and Alireza Tahbaz-Salehi. “Production Networks: A Primer”. In: *Annual Review of Economics* 11 (2019), pp. 635–663. DOI: 10.1146/annurev-economics-080218-030212.

[23] Plinio Centoamore and Luiz Fernando Rodrigues Pinto. “Remanufacturing Assessment of Machine Tools under a Circular Economy Perspective: A Resource Conservation Initiative”. In: *Sustainability* 16.8, 3109 (2024). DOI: 10.3390/su16083109.

[24] G. A. Cohen. *Karl Marx’s Theory of History: A Defence*. Princeton University Press, 1978.

[25] Renato Crouzeilles et al. “Ecological Restoration Success Is Higher for Natural Regeneration than for Active Restoration in Tropical Forests”. In: *Science Advances* 3.11 (2017), e1701345. DOI: 10.1126/sciadv.1701345.

[26] Will Damron. “Gains from factory electrification: Evidence from North Carolina, 1905–1926”. In: *Explorations in Economic History* 96 (2025), p. 101654. DOI: 10.1016/j.eeh.2025.101654.

[27] Paul A. David. “The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox”. In: *The American Economic Review* 80.2 (1990), pp. 355–361. URL: https://www.jstor.org/stable/2006600 (visited on 10/26/2026).

[28] Michel De Lara and Luc Doyen. *Sustainable Management of Natural Resources: Mathematical Models and Methods*. Berlin: Springer, 2008. DOI: 10.1007/978-3-540-79074-7.

[29] Thomas Dietz, Elinor Ostrom, and Paul C. Stern. “The Struggle to Govern the Commons”. In: *Science* 302.5652 (2003), pp. 1907–1912. DOI: 10.1126/science.1091015.

[30] Douglas Donnellan, Andy Lawrence, and Rose Weinschenk. *Annual Outage Analysis 2025*. Uptime Institute, May 2, 2025. URL: https://intelligence.uptimeinstitute.com/resource/annual-outage-analysis-2025 (visited on 10/26/2026).

[31] Ron Eglash. “An Introduction to Generative Justice”. In: *Teknokultura: Revista de Cultura Digital y Movimientos Sociales* 13.2 (2016), pp. 369–404. DOI: 10.5209/rev_TEKN.2016.v13.n2.52847.

[32] Friedrich Engels. *Letter to Joseph Bloch, 21–22 September 1890*. 1890. URL: https://www.marxists.org/archive/marx/works/1890/letters/90_09_21b.htm (visited on 10/26/2026). The correspondence is dated 21[-22] September 1890 in collected editions; the linked English online rendering identifies 21 September. Translation provenance varies among online versions.

[33] European Commission. *Commission Recommendation (EU) 2021/2279 on the use of Environmental Footprint methods to measure and communicate the life cycle environmental performance of products and organisations*. Dec. 15, 2021. URL: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32021H2279 (visited on 10/27/2026).

[34] European Commission. *Directive on repair of goods: Directive (EU) 2024/1799*. June 13, 2024. URL: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32024L1799 (visited on 10/26/2026).

[35] European Commission. *Commission Report on the Energy Performance of Data Centres in the European Union*. COM(2026) 500 final. European Commission, 2026. URL: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52026DC0500 (visited on 10/26/2026).

[36] European Parliament and Council of the European Union. *Regulation (EU) 2024/1781 establishing a framework for the setting of ecodesign requirements for sustainable products*. June 13, 2024. URL: https://eur-lex.europa.eu/eli/reg/2024/1781/oj (visited on 10/26/2026).

[37] Eurostat. *Over 12% of materials in the EU come from recycling*. Nov. 19, 2025. URL: https://ec.europa.eu/eurostat/web/products-eurostat-news/w/ddn-20251119-1 (visited on 10/26/2026).

[38] *Evaluation Criteria for Reliability-Centered Maintenance (RCM) Processes*. SAE International. 2024. DOI: 10.4271/JA1011_202411.

[39] Federal Trade Commission. *Partnerships Between Cloud Service Providers and AI Developers: Staff Report*. Federal Trade Commission, Jan. 2025. URL: https://search.ftc.gov/news-events/news/press-releases/2025/01/ftc-issues-staff-report-ai-partnerships-investments-study (visited on 10/26/2026).

[40] John Bellamy Foster. “Marx’s Theory of Metabolic Rift: Classical Foundations for Environmental Sociology”. In: *American Journal of Sociology* 105.2 (1999), pp. 366–405. DOI: 10.1086/210315.

[41] John Bellamy Foster. *Marx’s Ecology: Materialism and Nature*. Monthly Review Press, 2000. ISBN: 9781583670125.

[42] Étienne Fournier et al. “Human-cobot Collaboration’s Impact on Success, Time Completion, Errors, Workload, Gestures and Acceptability during an Assembly Task”. In: *Applied Ergonomics* 119 (2024), p. 104306. DOI: 10.1016/j.apergo.2024.104306.


DRAFT

[43] Nancy Fraser. “Behind Marx’s Hidden Abode: For an Expanded Conception of Capitalism”. In: *New Left Review* 86 (2014), pp. 55–72. DOI: 10.64590/kdf. (Visited on 10/26/2026).

[44] Robert A. Frosch. “Industrial Ecology: A Philosophical Introduction”. In: *Proceedings of the National Academy of Sciences* 89.3 (1992), pp. 800–803. DOI: 10.1073/pnas.89.3.800.

[45] Robert A. Frosch and Nicholas E. Gallopoulos. “Strategies for Manufacturing”. In: *Scientific American* 261.3 (1989), pp. 144–152. DOI: 10.1038/scientificamerican0989-144.

[46] George D. Gann et al. “International Principles and Standards for the Practice of Ecological Restoration: Third Edition”. In: *Restoration Ecology* 34 (2026), e70441. DOI: 10.1111/rec.70441.

[47] Gary Gereffi, John Humphrey, and Timothy Sturgeon. “The Governance of Global Value Chains”. In: *Review of International Political Economy* 12.1 (2005), pp. 78–104. DOI: 10.1080/09692290500049805.

[48] Krista J. Gile and Mark S. Handcock. “Respondent-Driven Sampling: An Assessment of Current Methodology”. In: *Sociological Methodology* 40.1 (2010), pp. 285–327. DOI: 10.1111/j.1467-9531.2010.01223.x.

[49] Google. *Power Usage Effectiveness: 2025 PUE Yearly Report*. 2025. URL: https://www.datacenters.google/efficiency/ (visited on 10/26/2026).

[50] Georg Graetz and Guy Michaels. “Robots at Work”. In: *Review of Economics and Statistics* 100.5 (2018), pp. 753–768. DOI: 10.1162/rest_a_00754.

[51] Sanford J. Grossman and Oliver D. Hart. “The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration”. In: *Journal of Political Economy* 94.4 (1986), pp. 691–719. DOI: 10.1086/261404.

[52] Henry Hansmann. *The Ownership of Enterprise*. Cambridge, MA: The Belknap Press of Harvard University Press, 1996. ISBN: 0674649702.

[53] Gavin Harper et al. “Recycling lithium-ion batteries from electric vehicles”. In: *Nature* 575 (2019), pp. 75–86. DOI: 10.1038/s41586-019-1682-5.

[54] Miguel A. Hernán and James M. Robins. “Using Big Data to Emulate a Target Trial When a Randomized Trial Is Not Available”. In: *American Journal of Epidemiology* 183.8 (2016), pp. 758–764. DOI: 10.1093/aje/kwv254.

[55] C. S. Holling. “Resilience and stability of ecological systems”. In: *Annual Review of Ecology and Systematics* 4 (1973), pp. 1–23. DOI: 10.1146/annurev.es.04.110173.000245.

[56] Michael G. Hudgens and M. Elizabeth Halloran. “Toward Causal Inference With Interference”. In: *Journal of the American Statistical Association* 103.482 (2008), pp. 832–842. DOI: 10.1198/016214508000000292.

[57] Institute of Electrical and Electronics Engineers and International Electrotechnical Commission. *IEC/IEEE International Standard: Time-Sensitive Networking Profile for Industrial Automation*. IEEE/IEC 60802-2026. June 29, 2026. URL: https://standards.ieee.org/ieee/60802/11358/ (visited on 10/26/2026).

[58] Intergovernmental Panel on Climate Change. *Climate Change 2022: Mitigation of Climate Change, Working Group III, Chapter 11: Industry*. IPCC, 2022. DOI: 10.1017/9781009157926.013. (Visited on 10/27/2026).

[59] Intergovernmental Science-Policy Platform on Biodiversity and Ecosystem Services. *Global Assessment Report on Biodiversity and Ecosystem Services*. IPBES, 2019. URL: https://www.ipbes.net/global-assessment (visited on 10/26/2026).

[60] Intergovernmental Science-Policy Platform on Biodiversity and Ecosystem Services. *Methodological Assessment Report on the Diverse Values and Valuation of Nature*. IPBES, 2022. URL: https://www.ipbes.net/the-values-assessment (visited on 10/27/2026).

[61] Intergovernmental Science-Policy Platform on Biodiversity and Ecosystem Services. *Thematic Assessment of the Interlinkages among Biodiversity, Water, Food and Health: Summary for Policymakers*. IPBES, 2024. URL: https://www.ipbes.net/nexus-assessment (visited on 10/26/2026).

[62] International Electrotechnical Commission. *IEC 62264-1:2013: Enterprise-Control System Integration—Part 1: Models and Terminology*. IEC 62264-1:2013. 2013. URL: https://webstore.iec.ch/en/publication/6675 (visited on 10/26/2026).

[63] International Electrotechnical Commission. *Failure Modes and Effects Analysis (FMEA and FMECA)*. IEC 60812:2018. 2018. URL: https://webstore.iec.ch/en/publication/26359 (visited on 10/26/2026).

[64] International Energy Agency. *Energy Efficiency 2025: Industry*. International Energy Agency, 2025. URL: https://www.iea.org/reports/energy-efficiency-2025/industry (visited on 10/26/2026).

[65] International Energy Agency. *Global Critical Minerals Outlook 2025*. Paris: International Energy Agency, 2025. URL: https://www.iea.org/reports/global-critical-minerals-outlook-2025 (visited on 10/26/2026).

[66] International Energy Agency. *Renewables for Industry*. International Energy Agency, Dec. 23, 2025. URL: https://www.iea.org/reports/renewables-for-industry (visited on 10/26/2026).

[67] International Energy Agency. *Global Critical Minerals Outlook 2026*. Paris: International Energy Agency, 2026. URL: https://www.iea.org/reports/global-critical-minerals-outlook-2026 (visited on 10/26/2026).

[68] International Energy Agency. *Key Questions on Energy and AI*. Paris: International Energy Agency, 2026. URL: https://www.iea.org/reports/key-questions-on-energy-and-ai (visited on 10/26/2026).

[69] International Federation of Robotics. *Five Million Robots now Operate in Factories Globally: World Robotics 2026 Report*. Sept. 24, 2026. URL: https://ifr.org/ifr-press-releases/five-million-robots-now-operate-in-factories-globally (visited on 10/26/2026).

[70] International Federation of Robotics. *President’s Report by Jane Heffner: World Robotics 2026*. Sept. 24, 2026. URL: https://ifr.org/ifr-press-releases/news/presidents-report-by-jane-heffner-3-2026 (visited on 10/26/2026).

[71] International Federation of Robotics. *Robot History*. URL: https://ifr.org/robot-history (visited on 10/26/2026).

[72] International Finance Corporation. *Good Practice Handbook on Cumulative Impact Assessment and Management: Guidance for the Private Sector in Emerging Markets*. Washington, DC: International Finance Corporation, 2013. URL: https://www.ifc.org/en/insights-reports/2013/publications-handbook-cumulativeimpactassessment (visited on 10/27/2026).

[73] International Labour Organization. *The Impact of Care Responsibilities on Women’s Labour Force Participation*. ILO Statistical Brief. International Labour Organization, Oct. 29, 2024. URL: https://www.ilo.org/sites/default/files/2024-10/GEDI-STAT%20brief_formatted_28.10.24_final.pdf (visited on 10/26/2026).

DRAFT

[74] International Labour Organization. *Generative AI and Jobs: A Refined Global Index of Occupational Exposure*. 140. International Labour Organization, May 20, 2025. DOI: 10.54394/HETP0387. (Visited on 10/26/2026).

[75] International Labour Organization. *Survey on Workers on Web-based Digital Platforms: Latin America and the Caribbean*. International Labour Organization, Apr. 10, 2025. DOI: 10.54394/ADTS8080.

[76] International Labour Organization. *Changing Landscape of Skills in the Age of AI*. International Labour Organization, Aug. 13, 2026. URL: https://www.ilo.org/publications/changing-landscape-skills-age-ai (visited on 10/26/2026).

[77] International Labour Organization. *Disruption without Dividend? How the Digital Divide and Task Differences Split GenAI's Global Impact*. ILO Working Paper 166. International Labour Organization, 2026. URL: https://www.ilo.org/publications/disruption-without-dividend-how-digital-divide-and-task-differences-split (visited on 10/26/2026).

[78] International Labour Organization. *Gen AI, Occupational Segregation and Gender Equality in the World of Work*. International Labour Organization, Mar. 5, 2026. DOI: 10.54394/00033798.

[79] International Labour Organization. *The World of Work and the 2030 Agenda: A Ten-year Review*. ILOSTAT, 2026. URL: https://ilostat.ilo.org/blog/the-world-of-work-and-the-2030-agenda-a-ten-year-review/ (visited on 10/26/2026).

[80] International Labour Organization. *Workers' Exposure to AI: What Indicators Tell Us and What They Don't*. International Labour Organization, Apr. 17, 2026. DOI: 10.54394/00033279. (Visited on 10/27/2026).

[81] International Organization for Standardization. *ISO 9283:1998: Manipulating Industrial Robots—Performance Criteria and Related Test Methods*. ISO 9283:1998. Apr. 1998. URL: https://www.iso.org/standard/22244.html (visited on 10/26/2026).

[82] International Organization for Standardization. *Condition Monitoring and Diagnostics of Machines: Data Processing, Communication and Presentation: Part 1, General Guidelines*. ISO 13374-1:2003. 2003. URL: https://www.iso.org/standard/21832.html (visited on 10/26/2026).

[83] International Organization for Standardization. *Environmental management — Life cycle assessment — Requirements and guidelines*. ISO 14044:2006. 2006. URL: https://www.iso.org/standard/38498.html (visited on 10/26/2026).

[84] International Organization for Standardization. *ISO 14040:2006: Environmental management — Life cycle assessment — Principles and framework*. 2006. URL: https://www.iso.org/standard/37456.html (visited on 10/26/2026).

[85] International Organization for Standardization. *Environmental management — Water footprint — Principles, requirements and guidelines*. ISO 14046:2014. 2014. URL: https://www.iso.org/standard/43263.html (visited on 10/27/2026).

[86] International Organization for Standardization. *ISO 22400-2:2014: Automation Systems and Integration — Key Performance Indicators (KPIs) for Manufacturing Operations Management — Part 2: Definitions and Descriptions*. 2014. URL: https://www.iso.org/standard/54497.html.

[87] International Organization for Standardization. *Condition Monitoring and Diagnostics of Machines: General Guidelines*. ISO 17359:2018. 2018. URL: https://www.iso.org/standard/71194.html (visited on 10/26/2026).

[88] International Organization for Standardization. *ISO 23247-1:2021: Automation Systems and Integration — Digital Twin Framework for Manufacturing — Part 1: Overview and General Principles*. 2021. URL: https://www.iso.org/standard/75066.html (visited on 10/27/2026).

[89] International Organization for Standardization. *ISO 23247-1:2021: Automation Systems and Integration—Digital Twin Framework for Manufacturing—Part 1: Overview and General Principles*. ISO 23247-1:2021. 2021. URL: https://www.iso.org/standard/75066.html (visited on 10/26/2026).

[90] International Organization for Standardization. *ISO 23247-2:2021: Automation Systems and Integration—Digital Twin Framework for Manufacturing—Part 2: Reference Architecture*. ISO 23247-2:2021. 2021. URL: https://www.iso.org/standard/78743.html (visited on 10/26/2026).

[91] International Organization for Standardization. *ISO 3691-4:2023: Industrial Trucks—Safety Requirements and Verification—Part 4: Driverless Industrial Trucks and Their Systems*. ISO 3691-4:2023. June 2023. URL: https://www.iso.org/standard/83545.html (visited on 10/26/2026).

[92] International Organization for Standardization. *Asset Management: Asset Management System: Requirements*. ISO 55001:2024. 2024. URL: https://www.iso.org/standard/83054.html (visited on 10/26/2026).

[93] International Organization for Standardization. *Circular economy — Measuring and assessing circularity performance*. ISO 59020:2024. 2024. URL: https://www.iso.org/standard/80650.html (visited on 10/26/2026).

[94] International Organization for Standardization. *ISO 10218-1:2025: Robotics—Safety Requirements—Part 1: Industrial Robots*. ISO 10218-1:2025. Feb. 2025. URL: https://www.iso.org/standard/73933.html (visited on 10/26/2026).

[95] International Organization for Standardization. *ISO 10218-2:2025: Robotics—Safety Requirements—Part 2: Industrial Robot Applications and Robot Cells*. ISO 10218-2:2025. Feb. 2025. URL: https://www.iso.org/standard/73934.html (visited on 10/26/2026).

[96] International Organization for Standardization and International Electrotechnical Commission. *ISO/IEC 17025:2017: General Requirements for the Competence of Testing and Calibration Laboratories*. ISO/IEC 17025:2017. 2017. URL: https://www.iso.org/standard/66912.html (visited on 10/26/2026).

[97] International Organization for Standardization and International Electrotechnical Commission. *Information Technology—Data Centres Key Performance Indicators—Part 9: Water Usage Effectiveness (WUE)*. ISO/IEC 30134-9:2022. 2022. URL: https://www.iso.org/standard/77692.html (visited on 10/26/2026).

[98] International Organization for Standardization and International Electrotechnical Commission. *Internet of Things (IoT) – Reference Architecture*. ISO/IEC 30141:2024. 2024. URL: https://www.iso.org/standard/88800.html (visited on 10/26/2026).

[99] International Society of Automation. *ISA-95 Series of Standards: Enterprise-Control System Integration*. URL: https://www.isa.org/standards-and-publications/isa-standards/isa-95-standard (visited on 10/26/2026).

[100] International Society of Automation. *ISA/IEC 62443 Series of Standards: Security for Industrial Automation and Control Systems*. URL: https://www.isa.org/standards-and-publications/isa-standards/isa-iec-62443-series-of-standards (visited on 10/26/2026).

[101] International Telecommunication Union and United Nations Institute for Training and Research. *The Global E-waste Monitor 2024*. ITU and UNITAR, 2024. URL: https://www.itu.int/en/ITU-D/Environment/Pages/Publications/The-Global-E-waste-Monitor-2024.aspx (visited on 10/26/2026).

DRAFT

[102] John P. A. Ioannidis. “Why Most Published Research Findings Are False”. In: *PLOS Medicine* 2.8 (2005), e124. DOI: 10.1371/journal.pmed.0020124. A 2022 correction (doi:10.1371/journal.pmed.1004085) restores missing parentheses in an equation in Table 2.

[103] *ISO 55000:2024: Asset management — Vocabulary, overview and principles*. International Organization for Standardization. 2024. URL: https://www.iso.org/standard/83053.html (visited on 10/26/2026).

[104] Andrew K. S. Jardine, Daming Lin, and Dragan Banjevic. “A Review on Machinery Diagnostics and Prognostics Implementing Condition-Based Maintenance”. In: *Mechanical Systems and Signal Processing* 20.7 (2006), pp. 1483–1510. DOI: 10.1016/j.ymssp.2005.09.012.

[105] Andrew Jarvis and Carey W. King. “Economy-wide Rebound and the Returns on Investment in Energy Efficiency”. In: *Energy Efficiency* 17, 60 (2024). DOI: 10.1007/s12053-024-10236-7. (Visited on 10/27/2026).

[106] Matthew Johnson and Jeffrey M. Bradshaw. “How Interdependence Explains the World of Teamwork”. In: *Engineering Artificially Intelligent Systems*. Vol. 13000. Lecture Notes in Computer Science. Springer, 2021, pp. 122–146. DOI: 10.1007/978-3-030-89385-9_8.

[107] Joint Committee for Guides in Metrology. *Evaluation of Measurement Data — Guide to the Expression of Uncertainty in Measurement*. JCGM 100:2008. 2008. DOI: 10.59161/JCGM100-2008E.

[108] Daniel O. B. Jones, Maria Belen Arias, Loïc Van Audenhaege, et al. “Long-Term Impact and Biological Recovery in a Deep-Sea Mining Track”. In: *Nature* 642 (2025), pp. 112–118. DOI: 10.1038/s41586-025-08921-3.

[109] Slimane Kerrouchi, El-Houssaine Aghezzaf, and Johannes Cottyn. “Production digital twin: a systematic literature review of challenges”. In: *International Journal of Computer Integrated Manufacturing* 37.10-11 (2024), pp. 1168–1193. DOI: 10.1080/0951192X.2024.2314792.

[110] Yaguo Lei et al. “Machinery Health Prognostics: A Systematic Review from Data Acquisition to RUL Prediction”. In: *Mechanical Systems and Signal Processing* 104 (2018), pp. 799–834. DOI: 10.1016/j.ymssp.2017.11.016.

[111] Manfred Lenzen. “Errors in Conventional and Input–Output-based Life-cycle Inventories”. In: *Journal of Industrial Ecology* 4.4 (2000), pp. 127–148. DOI: 10.1162/10881980052541981.

[112] Wassily W. Leontief. “Quantitative Input and Output Relations in the Economic Systems of the United States”. In: *The Review of Economics and Statistics* 18.3 (1936), pp. 105–125. DOI: 10.2307/1927837.

[113] Library of Congress. *Ford Implements the Moving Assembly Line*. Library of Congress. URL: https://guides.loc.gov/this-month-in-business-history/october/ford (visited on 10/26/2026).

[114] Roderick J. A. Little and Donald B. Rubin. *Statistical Analysis with Missing Data*. 3rd ed. Wiley, 2019. DOI: 10.1002/9781119482260.

[115] Mei Lu and Ke Wang. “The Direct Energy Rebound Effects for Manufacturing and Service Sectors in China: Evidence from Firm-Level Estimations”. In: *Journal of Cleaner Production* 474, 143617 (2024). DOI: 10.1016/j.jclepro.2024.143617. (Visited on 10/27/2026).

[116] Yan Lu et al. *An Overarching Quality Evaluation Framework for Additive Manufacturing Digital Twin*. National Institute of Standards and Technology, 2024. URL: https://www.nist.gov/publications/overarching-quality-evaluation-framework-additive-manufacturing-digital-twin (visited on 10/26/2026).

[117] Benjamin P. MacLeod et al. “Self-Driving Laboratory for Accelerated Discovery of Thin-Film Materials”. In: *Science Advances* 6.20, eaaz8867 (2020). DOI: 10.1126/sciadv.aaz8867.

[118] Charles F. Manski. *Partial Identification of Probability Distributions*. Springer, 2003. DOI: 10.1007/b97478.

[119] Karl Marx. *Preface to A Contribution to the Critique of Political Economy*. 1859. URL: https://www.marxists.org/archive/marx/works/1859/critique-pol-economy/preface.htm (visited on 10/26/2026). The online English text is based on the Progress Publishers translation, Moscow, 1977; original work published 1859.

[120] Karl Marx. *Capital, Volume I, Chapter 15: Machinery and Modern Industry*. 1867. URL: https://www.marxists.org/archive/marx/works/1867-c1/ch15.htm (visited on 10/26/2026). German original 1867; online English edition after the Moore–Aveling translation (1887), edited by Engels; chapter numbering edition-specific.

[121] Karl Marx. *Capital, Volume I, Chapter 23: Simple Reproduction*. 1867. URL: https://www.marxists.org/archive/marx/works/1867-c1/ch23.htm (visited on 10/26/2026). German original 1867; online English edition after the Moore–Aveling translation (1887), edited by Engels; chapter numbering edition-specific.

[122] Karl Marx. *Capital, Volume I, Chapter 24: Conversion of Surplus-Value into Capital*. 1867. URL: https://www.marxists.org/archive/marx/works/1867-c1/ch24.htm (visited on 10/26/2026). German original 1867; online English edition after the Moore–Aveling translation (1887), edited by Engels; chapter numbering edition-specific.

[123] Karl Marx. *Capital, Volume I, Chapter 7: The Labour-Process and the Process of Producing Surplus-Value*. 1867. URL: https://www.marxists.org/archive/marx/works/1867-c1/ch07.htm (visited on 10/26/2026). German original 1867; online English edition after the Moore–Aveling translation (1887), edited by Engels; chapter numbering edition-specific.

[124] Karl Marx. *Capital, Volume II, Chapter 20: Simple Reproduction*. 1885. URL: https://www.marxists.org/archive/marx/works/1885-c2/ch20_01.htm (visited on 10/26/2026). Originally published in 1885 after editorial preparation by Friedrich Engels. The linked English text follows the revised 1893 German edition as translated for Progress Publishers, 1956.

[125] Karl Marx. *Grundrisse: Notebook VII, Machines and Fixed Capital (online translated edition)*. Trans. by Martin Nicolaus. 1973. URL: https://www.marxists.org/archive/marx/works/1857/grundrisse/ch14.htm (visited on 10/26/2026). Drafted 1857–1858; cited English translation (Nicolaus 1973); German editions 1939–1941 and 1953.

[126] David Q. Mayne et al. “Constrained Model Predictive Control: Stability and Optimality”. In: *Automatica* 36.6 (2000), pp. 789–814. DOI: 10.1016/S0005-1098(99)00214-9.

[127] MIT Laboratory for Information and Decision Systems. *History*. URL: https://www.lids.mit.edu/about-lids/history (visited on 10/26/2026).

[128] MIT Laboratory for Information and Decision Systems. *The Man Who Came to Dinner: Servomechanisms Laboratory History and Numerical Control*. URL: https://lidsmag.lids.mit.edu/2006/the_man_who_came_to_dinner.html (visited on 10/26/2026).

DRAFT

[129] David Moreno-Mateos et al. “Structural and Functional Loss in Restored Wetland Ecosystems”. In: *PLoS Biology* 10.1 (2012), e1001247. DOI: 10.1371/journal.pbio.1001247.

[130] National Aeronautics and Space Administration. *Prognostics Center of Excellence Data Set Repository*. URL: https://www.nasa.gov/intelligent-systems-division/discovery-and-systems-health/pcoe/pcoe-data-set-repository/ (visited on 10/26/2026).

[131] National Center for Science and Engineering Statistics. *The State of U.S. Science and Engineering 2026*. National Science Foundation. 2026. URL: https://ncses.nsf.gov/pubs/nsbsep20261 (visited on 10/26/2026).

[132] National Institute of Standards and Technology. *Performance Assessment Framework for Robotic Systems*. Mar. 26, 2025. URL: https://www.nist.gov/programs-projects/performance-assessment-framework-robotic-systems (visited on 10/26/2026).

[133] National Institute of Standards and Technology. *Digital Thread for Smart Manufacturing*. URL: https://www.nist.gov/programs-projects/digital-thread-smart-manufacturing (visited on 10/26/2026).

[134] National Laboratory of the Rockies. *Black Start: Grid Modernization*. URL: https://www.nlr.gov/grid/black-start (visited on 10/26/2026).

[135] David F. Noble. *Forces of Production: A Social History of Industrial Automation*. New York: Alfred A. Knopf, 1984. URL: https://www.si.edu/object/forces-production-social-history-industrial-automation-david-f-noble%3Asiris_sil_204376 (visited on 10/26/2026).

[136] David E. Nye. “What Was the Assembly Line?” In: *Temp: Tidsskrift for Historie* 1.1 (2010), pp. 59–81. URL: https://tidsskrift.dk/temp/article/view/24974 (visited on 10/26/2026).

[137] Thomas O’Neill et al. “Human–Autonomy Teaming: A Review and Analysis of the Empirical Literature”. In: *Human Factors* 64.5 (2022). DOI: 10.1177/0018720820960865.

[138] Occupational Safety and Health Administration. *29 CFR 1910.147: The Control of Hazardous Energy (Lockout/Tagout)*. URL: https://www.osha.gov/laws-regs/regulations/standardnumber/1910/1910.147 (visited on 10/26/2026).

[139] OECD. *Guide to OECD Trade in Value Added (TiVA) Indicators, 2025 Edition*. OECD, Directorate for Science, Technology and Innovation, Sept. 2025. URL: https://stats.oecd.org/wbos/fileview2.aspx?IDFile=2143f34e-6feb-41a9-abaf-cb52132608c4 (visited on 10/26/2026).

[140] OECD. *How Widespread Is Algorithmic Management in Workplaces?* Paris: OECD Publishing, Dec. 19, 2025. DOI: 10.1787/cda7a114-en.

[141] OECD. *Mapping the Semiconductor Value Chain: Working towards Identifying Dependencies and Vulnerabilities*. OECD Science, Technology and Industry Policy Papers 182. Paris, 2025. DOI: 10.1787/4154cdbf-en.

[142] OECD. *OECD Supply Chain Resilience Review: Navigating Risks*. Paris: OECD, June 2025. DOI: 10.1787/94e3a8ea-en. (Visited on 10/26/2026).

[143] OECD. *The Chip Landscape: Geographical Distribution of Wafer Fabrication Capacity*. OECD Science, Technology and Industry Policy Papers 188. Paris, Dec. 1, 2025. DOI: 10.1787/02dbd028-en.

[144] OECD. *Trade in Value Added (TiVA) 2025 Edition: Principal Indicators*. Jan. 2026. URL: https://data-explorer.oecd.org/ (visited on 10/26/2026).

[145] Open X-Embodiment Collaboration. “Open X-Embodiment: Robotic Learning Datasets and RT-X Models”. In: *Proceedings of the IEEE International Conference on Robotics and Automation*. 2024, pp. 6892–6903. DOI: 10.1109/ICRA57147.2024.10611477. arXiv: 2310.08864. (Visited on 10/26/2026).

[146] Elinor Ostrom. *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge University Press, 1990. DOI: 10.1017/CBO9780511807763.

[147] Elinor Ostrom. “A General Framework for Analyzing Sustainability of Social-Ecological Systems”. In: *Science* 325.5939 (2009), pp. 419–422. DOI: 10.1126/science.1172133.

[148] Elinor Ostrom. “Beyond Markets and States: Polycentric Governance of Complex Economic Systems”. In: *American Economic Review* 100.3 (2010), pp. 641–672. DOI: 10.1257/aer.100.3.641.

[149] Raja Parasuraman, Thomas B. Sheridan, and Christopher D. Wickens. “A Model for Types and Levels of Human Interaction with Automation”. In: *IEEE Transactions on Systems, Man, and Cybernetics—Part A: Systems and Humans* 30.3 (2000), pp. 286–297. DOI: 10.1109/3468.844354.

[150] Matteo Pasquinelli. *The Eye of the Master: A Social History of Artificial Intelligence*. London: Verso, 2023. ISBN: 9781788730068. URL: https://www.versobooks.com/products/735-the-eye-of-the-master (visited on 10/26/2026).

[151] Virginie Pérotin. “What Do We Really Know about Workers’ Co-operatives?” In: *Mainstreaming Co-operation: An Alternative for the Twenty-First Century?* Ed. by Anthony Webster, Linda Shaw, and Rachael Vorberg-Rugh. Manchester University Press, 2016, pp. 239–260. DOI: 10.7228/manchester/9780719099595.003.0014.

[152] Jonas Peters, Dominik Janzing, and Bernhard Sch”olkopf. *Elements of Causal Inference: Foundations and Learning Algorithms*. MIT Press, 2017. ISBN: 9780262037310.

[153] Jeffrey Pfeffer and Gerald R. Salancik. *The External Control of Organizations: A Resource Dependence Perspective*. New York: Harper & Row, 1978.

[154] Planetary Boundaries Science Lab. *Planetary Health Check 2026: A Scientific Assessment of the State of the Planet*. Potsdam: Potsdam Institute for Climate Impact Research, Sept. 21, 2026. DOI: 10.48485/pik.2026.25. (Visited on 10/26/2026).

[155] Richard J. Plevin, Mark A. Delucchi, and Felix Creutzig. “Using Attributional Life Cycle Assessment to Estimate Climate-Change Mitigation Benefits Misleads Policy Makers”. In: *Journal of Industrial Ecology* 18.1 (2014), pp. 73–83. DOI: 10.1111/jiec.12074.

[156] Antonio Possolo, Sally S. Bruce, and Jr. Watters Robert L. *Metrological Traceability: Frequently Asked Questions and NIST Policy*. NIST Technical Note 2156. National Institute of Standards and Technology, 2021. DOI: 10.6028/NIST.TN.2156. (Visited on 10/27/2026).

[157] Guixiu Qiao and Brian A. Weiss. *Accuracy Degradation Analysis for Industrial Robot Systems*. National Institute of Standards and Technology, June 7, 2017. URL: https://www.nist.gov/publications/accuracy-degradation-analysis-industrial-robot-systems (visited on 10/26/2026).

DRAFT

[158] James B. Rawlings, David Q. Mayne, and Moritz M. Diehl. *Model Predictive Control: Theory, Computation, and Design*. 2nd ed. Nob Hill Publishing, 2026. URL: https://sites.chemengr.ucsb.edu/~jbraw/mpc/ (visited on 10/27/2026).

[159] Johan Rockstr"om, Joyeeta Gupta, Dahe Qin, et al. "Safe and just Earth system boundaries". In: *Nature* 619 (2023), pp. 102–111. DOI: 10.1038/s41586-023-06083-8.

[160] Paul R. Rosenbaum. *Observational Studies*. 2nd ed. Springer, 2002. DOI: 10.1007/978-1-4757-3692-2.

[161] Paul R. Rosenbaum and Donald B. Rubin. "The Central Role of the Propensity Score in Observational Studies for Causal Effects". In: *Biometrika* 70.1 (1983), pp. 41–55. DOI: 10.1093/biomet/70.1.41.

[162] Kohei Saito. *Karl Marx's Ecosocialism: Capital, Nature, and the Unfinished Critique of Political Economy*. Monthly Review Press, 2017. ISBN: 9781583676400.

[163] Nithya Sambasivan et al. ""Everyone wants to do the model work, not the data work": Data Cascades in High-Stakes AI". In: *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, 2021. DOI: 10.1145/3411764.3445518.

[164] Mahadev Satyanarayanan. "The Emergence of Edge Computing". In: *Computer* 50.1 (2017), pp. 30–39. DOI: 10.1109/MC.2017.9.

[165] Schneider Electric. *Modicon: History of Programmable Logic Controllers*. URL: https://www.se.com/uk/en/about-us/company-profile/brands/modicon/ (visited on 10/26/2026).

[166] Marvin Schütt et al. "How to Explain the Huge Differences in Rebound Estimates: A Meta-Regression Analysis of the Literature". In: *The Energy Journal* 46.2 (2025). DOI: 10.1177/01956574241289004. (Visited on 10/27/2026).

[167] SEMI. *Global Semiconductor Equipment Sales Forecast to Reach a Record $229 Billion in 2028, SEMI Reports*. July 14, 2026. URL: https://www.semi.org/en/semi-press-release/global-semiconductor-equipment-sales-forecast-to-reach-a-record-229-billion-dollars-in-2028-semi-reports (visited on 10/26/2026).

[168] Guodong Shao, Deogratias Kibira, and Simon P. Frechette. *Digital Twins for Advanced Manufacturing: The Standardized Approach*. National Institute of Standards and Technology, 2024. URL: https://www.nist.gov/publications/digital-twins-advanced-manufacturing-standardized-approach (visited on 10/26/2026).

[169] Shuning Shi and Xiaoyu Yan. "A critical review on spatially explicit life cycle assessment methodologies and applications". In: *Sustainable Production and Consumption* 52 (2024), pp. 566–579. DOI: 10.1016/j.spc.2024.11.015.

[170] Chris Smith and Paul Thompson. "Braverman and Labor and Monopoly Capital: A Retrospective". In: *Contributions to Political Economy* 43.1 (2024), pp. 61–81. DOI: 10.1093/cpe/bzae007.

[171] Sarah J. Smith et al. *United States Data Center Energy Usage Report: 2025 Update*. Lawrence Berkeley National Laboratory, June 18, 2026. DOI: 10.71468/P1RP4F.

[172] Piero Sraffa. *Production of Commodities by Means of Commodities: Prelude to a Critique of Economic Theory*. Cambridge: Cambridge University Press, 1960.

[173] Konstantin Stadler et al. "EXIOBASE 3: Developing a Time Series of Detailed Environmentally Extended Multi-Regional Input–Output Tables". In: *Journal of Industrial Ecology* 22.3 (2018), pp. 502–515. DOI: 10.1111/jiec.12715.

[174] Stanford Institute for Human-Centered Artificial Intelligence. *Artificial Intelligence Index Report 2026*. Stanford University, 2026. URL: https://hai.stanford.edu/ai-index/2026-ai-index-report (visited on 10/26/2026).

[175] Sebastian Steiner et al. "Organic Synthesis in a Modular Robotic System Driven by a Chemical Programming Language". In: *Science* 363.6423, eaav2211 (2019). DOI: 10.1126/science.aav2211.

[176] Keith Stouffer et al. *Guide to Operational Technology (OT) Security*. NIST SP 800-82 Revision 3. National Institute of Standards and Technology, 2023. DOI: 10.6028/NIST.SP.800-82r3.

[177] Elizabeth A. Stuart et al. "The Use of Propensity Scores to Assess the Generalizability of Results from Randomized Trials". In: *Journal of the Royal Statistical Society: Series A (Statistics in Society)* 174.2 (2011), pp. 369–386. DOI: 10.1111/j.1467-985X.2010.00673.x.

[178] Nathan J. Szymanski et al. "An autonomous laboratory for the accelerated synthesis of inorganic materials". In: *Nature* 624 (2023), pp. 86–91. DOI: 10.1038/s41586-023-06734-w. The article was corrected on 19 January 2026; see the separate author-correction reference.

[179] Nathan J. Szymanski et al. "Author Correction: An autonomous laboratory for the accelerated synthesis of inorganic materials". In: *Nature* 650 (2026), E1. DOI: 10.1038/s41586-025-09992-y.

[180] Elham Tabassi. *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*. NIST AI 100-1. National Institute of Standards and Technology, 2023. DOI: 10.6028/NIST.AI.100-1.

[181] The Henry Ford. *Robot, First Unimate Robot Ever Installed on an Assembly Line, 1961*. URL: https://www.thehenryford.org/collections/explore/artifact/183434 (visited on 10/26/2026).

[182] U.S. Geological Survey. *Mineral Commodity Summaries 2026*. U.S. Geological Survey Mineral Commodity Summaries. May 2026. DOI: 10.3133/mcs2026. (Visited on 10/26/2026).

[183] United Nations Environment Programme. *A Policymakers' Guide to Life Cycle Assessment*. UNEP Life Cycle Initiative, Nov. 2024. URL: https://wedocs.unep.org/handle/20.500.11822/46469 (visited on 10/27/2026).

[184] United Nations Environment Programme. *Global Resources Outlook 2024: Bend the Trend: Pathways to a Liveable Planet as Resource Use Spikes*. International Resource Panel, United Nations Environment Programme, 2024. URL: https://www.unep.org/resources/Global-Resource-Outlook-2024 (visited on 10/26/2026).

[185] United States Department of Energy. *Cooling Water Efficiency Opportunities for Federal Data Centers*. Jan. 9, 2019. URL: https://www.energy.gov/cmei/femp/cooling-water-efficiency-opportunities-federal-data-centers (visited on 10/26/2026).

[186] USEtox International Center. *The USEtox Model: Consensus-based Characterization of Human Toxicity and Freshwater Ecotoxicity*. URL: https://usetox.org/model (visited on 10/27/2026).

[187] Otto Van Geet and David Sickinger. *Best Practices Guide for Energy-Efficient Data Center Design*. National Renewable Energy Laboratory and Federal Energy Management Program, United States Department of Energy, 2024. DOI: 10.2172/2417618. (Visited on 10/26/2026).

DRAFT

[188] Amanda A. Volk and Milad Abolhasani. “Performance Metrics to Unleash the Power of Self-Driving Labs in Chemistry and Materials Science”. In: *Nature Communications* 15, 1378 (2024). DOI: 10.1038/s41467-024-45569-5.

[189] Li-Chih Wang and Shao-Ying Wu. “Modeling with Colored Timed Object-Oriented Petri Nets for Automated Manufacturing Systems”. In: *Computers & Industrial Engineering* 34.2 (1998), pp. 463–480. DOI: 10.1016/S0360-8352(97)00145-9.

[190] Thomas Wiedmann and Manfred Lenzen. “Environmental and Social Footprints of International Trade”. In: *Nature Geoscience* 11 (2018), pp. 314–321. DOI: 10.1038/s41561-018-0113-9.

[191] Mark D. Wilkinson et al. “The FAIR Guiding Principles for scientific data management and stewardship”. In: *Scientific Data* 3 (2016), p. 160018. DOI: 10.1038/sdata.2016.18.

[192] Ellen Meiksins Wood. *The Origin of Capitalism: A Longer View*. Verso, 2002. ISBN: 9781859843925.

[193] World Bank. *World Development Report 2020: Trading for Development in the Age of Global Value Chains*. Washington, DC: World Bank, 2020. URL: https://www.worldbank.org/en/publication/wdr2020 (visited on 10/26/2026).

[194] World Intellectual Property Organization. *SPARK: Patent Trends Update in GenAI*. July 14, 2026. URL: https://www.wipo.int/web-publications/spark-patent-trends-update-in-genai/en/global-patenting-trends.html (visited on 10/26/2026).

[195] World Resources Institute and World Business Council for Sustainable Development. *Product Life Cycle Accounting and Reporting Standard*. Greenhouse Gas Protocol, 2011. URL: https://ghgprotocol.org/product-standard (visited on 10/27/2026).

[196] World Wide Web Consortium. *PROV-O: The PROV Ontology*. Apr. 30, 2013. URL: https://www.w3.org/TR/prov-o/ (visited on 10/26/2026).

DRAFT