Governance Near Criticality: Rethinking Intervention under Critical Transition Uncertainty

Wanhong HUANG
huangwanhong@serendip.ngo

Abstract

This paper develops a conceptual framework for political response near critical states. The point of departure is a limited but important implication of nonlinear dynamics: when a system approaches a critical transition, relatively small perturbations can sometimes produce changes that are disproportionate to their initial magnitude, while subsequent return may become difficult under hysteresis, path dependence, or self-reinforcing feedback. The political significance of this condition cannot be reduced to a problem of improved control. Near criticality, information about leverage can increase the ethical weight of intervention, because intentional steering may alter not only a present state but also the set of future trajectories available to affected actors. The paper therefore places dynamical analysis alongside epistemology, justice, ethics, aesthetics, and political legitimacy. Particular attention is given to random noise as a dialectical object. Fluctuations can obscure structure, reveal susceptibility, trigger transitions, or participate in the formation of new organization; their political meaning is consequently treated as relational and historically situated. The paper introduces a response-centered vocabulary in which suppression is one possible mode among observation, buffering, absorption, redirection, selective amplification, and restraint. Daoist ideas of *wuwei* and the imagery of water are used cautiously as comparative resources for thinking about non-coercive response without identifying political systems with natural flows. The analysis further examines the limits of tangent-space reasoning, the performative character of near-critical information in reflexive social systems, the justice of trajectory closure, and the value of reversibility and future option space. The paper advances a provisional orientation toward responsive participation rather than a universal prescription for near-critical governance.

Keywords: criticality; political dynamics; perturbation; noise; response; governance; irreversibility; justice; wuwei; option space

Discussion Paper Note

This paper examines political response in conditions provisionally described as *near criticality*. Its structural role is to connect three domains that are often treated separately: the dynamical study of critical transitions, the epistemology of decision under nonlinear sensitivity, and the normative study of political intervention. The paper treats their intersection as a problem of political judgment rather than as a ready-made application of control theory.

The term *near criticality* is used cautiously. In nonlinear dynamical systems, critical transitions can involve abrupt movement between contrasting regimes, and several families of models exhibit rising susceptibility, critical slowing down, noise-induced transitions, or hysteresis near thresholds (Scheffer, Carpenter, et al. 2001; Scheffer, Bascompte, et al. 2009; Forgoston and Moore 2018). These results motivate the present inquiry, but they do not establish that any particular political system is literally governed by the same equations. The paper therefore uses concepts such as attractor, basin, perturbation, phase transition, and critical region as formally suggestive terms whose political content requires independent theoretical and empirical justification.

The inquiry begins from a practical tension. Near a sensitive transition, small interventions may carry unusual leverage. That possibility can make control appear attractive: if a minor perturbation can help move a system away from an undesirable trajectory, intervention may seem efficient. Yet the same sensitivity raises the stakes of error. A locally small action can participate in a large and persistent transition, and the resulting trajectory can alter the future choice set of actors who did not authorize the intervention. The paper therefore separates *dynamical leverage* from *normative warrant*. Information that a system is susceptible to steering does not by itself establish that steering is justified.

A second focus concerns noise. Engineering and policy vocabularies often treat noise as a disturbance to be reduced so that a stable signal can be recovered. Nonlinear dynamics provides a more differentiated picture. Noise can induce escapes from quasi-stable states, and under some conditions it can contribute constructively to signal detection or organized response (Astumian and Moss 1998; Forgoston and Moore 2018). These physical results are not translated directly into political prescriptions. They motivate an epistemic question: before a fluctuation is classified as a defect to be removed, what does its interaction with the system reveal about susceptibility, coupling, hidden structure, and emerging organization?

The phrase *epistem the noise* is used in this exploratory sense. It asks that the political meaning of a fluctuation remain open long enough for its relational role to be investigated. A disturbance can obscure information, reveal an unstable coupling, release accumulated tension, initiate a damaging cascade, or contribute to a new pattern of coordination. The relevant political object is therefore not the perturbation in isolation but the perturbation–response relation. The paper develops this relation without assuming that its value can be inferred from dynamical form alone.

This position creates distance from a purely technocratic reading of adaptive governance. Existing work on complexity and governance already emphasizes nonlinearity, threshold effects, cascades, limited predictability, and the need for adaptive capacities (Duit and Galaz 2008; May 2022). The present inquiry extends that concern toward the normative consequences of high susceptibility. A policy actor near criticality may possess information about a locally influential direction while remaining uncertain about the justice of the resulting transition, the distribution of risk, the durability of the new regime, and the legitimacy of intentionally closing alternative future trajectories.

The paper consequently treats intervention and restraint as members of a wider response repertoire. Observation, buffering, dissipation, temporary decoupling, redirection, selective amplification, and intentional non-forcing can each become relevant under different configurations. The paper does not rank these responses in advance. Their appropriateness depends on the system's structure, the quality of available evidence, the distribution of expected burdens and benefits, and the extent to which later correction remains possible.

Daoist ideas of *wuwei* and the recurring image of water in the *Laozi* provide a comparative philosoph-

ical resource for this problem. Contemporary scholarship commonly cautions that *wuwei* should not be reduced to literal inactivity; interpretations such as non-assertive or non-coercive action capture part of its practical orientation, while water is used in the text to express yieldingness and strength (Chan 2025). This paper does not identify Daoist philosophy with complexity science. The comparison is narrower: it helps articulate a mode of response that can remain active without presuming that every obstacle should be overcome by direct force or that every future trajectory should be fixed in advance.

The normative dimension remains explicit throughout. Political intervention redistributes benefits, burdens, risks, opportunities, and future options. Distributive-justice theory provides multiple competing accounts of how such distributions should be evaluated (Lamont and Favor 2017), while moral decision under uncertainty raises further problems when consequences, probabilities, or even relevant moral principles remain uncertain (Tarsney et al. 2024). Near criticality intensifies these questions because a small action may be associated with a large change in trajectory. The paper therefore resists any inference from dynamical efficiency to moral legitimacy.

Aesthetic judgment enters the inquiry in a different role. Near a transition, several future configurations can remain partially manifest rather than fully specified. Political actors may perceive coherence, fragility, or emergent form before the normative consequences of that form are settled. Such perception can guide attention, but the paper does not treat aesthetic salience as a substitute for justice. A form can appear elegant or coherent while imposing serious burdens on others. The relation among aesthetics, ethics, and justice is therefore treated as one of differentiation and possible interaction rather than reduction.

The paper also develops a limitation of local, tangent-space reasoning. Local information can remain informative while the intervention justified by that information becomes difficult to contain. Near an attractor boundary or bifurcation, a small movement can alter the long-run regime. In reflexive social systems, publication of near-critical information can itself change behavior, expectations, and strategic interaction. The paper therefore distinguishes an epistemically supported neighborhood from an interventionally safe neighborhood. Their relation is treated as an open empirical and normative problem.

The intended contribution is consequently modest in form but broad in scope. The paper proposes a vocabulary for thinking about political response when susceptibility is elevated, prediction remains incomplete, and action can alter the future conditions of action. It advances neither a universal preference for intervention nor a universal preference for restraint. Its provisional orientation is toward responsive participation: observe the perturbation, characterize its relation to the system, protect against severe and persistent harm where possible, preserve revisability where feasible, and subject any intentional trajectory-shaping to explicit questions of justice and legitimacy.

Several research obligations remain. The first concerns empirical identification: political systems may display transition-like behavior without satisfying the mathematical conditions associated with a specific critical phenomenon. The second concerns measurement: susceptibility, noise structure, network coupling, and basin geometry are rarely directly observable in social systems. The third concerns normativity: concepts such as option preservation, reversibility, and non-coercive response require justification when they conflict with urgent protection, distributive justice, or democratic decision. The fourth concerns historical variation: forms of legitimate response may differ across institutions, scales, and periods. These obligations are treated as reasons for continued inquiry rather than as defects to be concealed by a stronger vocabulary.

Responsible Use and Rights Reservation

This section separates requested scholarly conduct from the legal permissions stated on the following page. It records an ethical request for responsible use and then defines the narrower scope of retained legal rights.

The author encourages good-faith discussion, criticism, independent inquiry, and responsible use of the material in this work. The framework developed here concerns conditions in which political interventions may have disproportionate or persistent consequences. Readers are therefore asked to treat its concepts as objects for critical inquiry rather than as instructions for manipulating real political systems. This ethical request leaves the licence's permissions and legally authorized uses unchanged.

The author retains the rights preserved under CC BY-NC 4.0 and may pursue remedies to which the author is legally entitled for breach of the licence or violation of independently applicable rights. Reuse remains independent from authorial endorsement. Third-party rights require authorization from their respective holders where applicable. Copyright exceptions and limitations, including applicable forms of fair use or fair dealing, remain fully available.

Notices

Status. This working draft records an early stage of a conceptual inquiry. The proposed vocabulary, formal sketches, analogies, and normative distinctions remain open to revision. The paper does not assume that concepts developed in statistical physics, ecology, or stochastic dynamical systems transfer directly to international or political systems. Empirical identification of near-critical political configurations, measurement of susceptibility, and estimation of transition probabilities remain separate research obligations.

Statement on the use of language models. The exploratory discussions and preparation of this paper involved OpenAI's ChatGPT. ChatGPT supported exploratory dialogue, source discovery followed by website verification, conceptual reconstruction, argumentative criticism, and LaTeX drafting. The author selected the research problem, directed and approved theoretical commitments, and bears sole responsibility for the manuscript, including its definitions, claims, interpretations, and errors. Authorship credit remains with the human author.

1 Introduction

This section establishes the structural problem of the paper. Its objective is to distinguish near-critical political response from ordinary state-targeting and to clarify why increased dynamical leverage can simultaneously enlarge the ethical and epistemic burden of intervention. The discussion proceeds from critical transitions and perturbations to response, normativity, and the limits of steering. It then states the paper's working concepts and the research obligations required before these concepts can support empirical political analysis.

Complex systems can display transitions in which gradual changes in conditions are followed by comparatively abrupt movement into a different regime. Research across ecology, climate, and other dynamical domains has examined tipping behavior, alternative stable states, loss of resilience, and early-warning patterns associated with some classes of transition (Scheffer, Carpenter, et al. 2001; Scheffer, Bascompte, et al. 2009). Stochastic dynamics adds a further complication: irregular perturbations can induce escapes from quasi-stable states even when deterministic forcing alone would not yet produce a transition (Forgoston and Moore 2018). These results make criticality attractive as a language for politics because they offer a way to think about periods in which ordinary proportional intuitions become unreliable. Their political use nevertheless requires restraint. The presence of abrupt political change, polarization, institutional fragility, or strategic uncertainty does not by itself demonstrate a mathematical critical point.

The paper therefore begins with a provisional concept rather than an empirical classification.

Definition 1.1 (Near-critical political configuration)**.** A *near-critical political configuration* is a hypothesized condition in which the reachable political trajectory is unusually sensitive to a subset of perturbations, while multiple materially different future regimes remain plausibly accessible. The definition is relational and diagnostic. It does not presume a particular mathematical mechanism, a universal early-warning signal, or a uniquely desirable successor regime.

This definition shifts attention from the observed state alone toward the way a system responds. Two political configurations can appear similar in aggregate indicators while differing considerably in susceptibility. Let x_t denote a coarse representation of the current configuration, G_t a changing coupling structure, and η_t a perturbation. A minimal response representation is

$$x_{t+1} = F(x_t, G_t, \eta_t, u_t), \tag{1}$$

where u_t denotes intentional intervention. Equation (1) is not offered as an estimable model of international politics. It marks a conceptual distinction: the effect of a perturbation depends on the state and relational architecture into which it enters, while the intervention itself enters the same evolving system.

Figure 1 shows the intuition that motivates the paper. If a configuration lies deep inside a relatively stable basin, a small perturbation may dissipate without altering the long-run regime. Near a boundary, a perturbation of similar magnitude may contribute to a different trajectory. The same possibility that creates leverage also creates ethical weight. A political actor who can influence a highly sensitive system may be able to obtain a large effect with a small intervention, while affected actors bear consequences whose scale greatly exceeds the initiating act.

This asymmetry separates *dynamical efficiency* from *normative legitimacy*. Low intervention cost says little about the justice of the resulting transition. A movement toward a new regime can redistribute security, wealth, status, political standing, environmental burden, and future agency. Distributive-justice theory offers several competing principles for evaluating such distributions, and their disagreement is itself relevant to political choice (Lamont and Favor 2017). Under uncertainty, the difficulty becomes sharper because actors may remain unsure about both consequences and normative principles (Tarsney

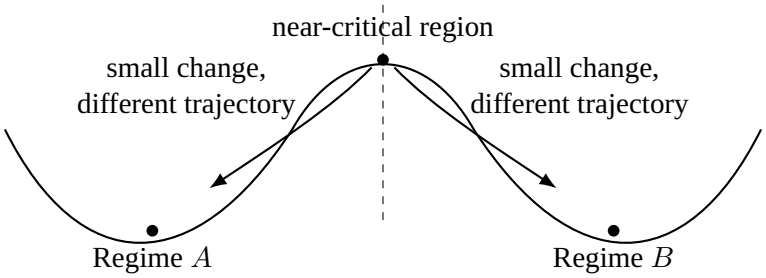

Figure 1: Schematic relation between a near-critical region and contrasting long-run trajectories. The figure illustrates sensitivity and possible trajectory divergence rather than an empirical claim about a particular political system.

et al. 2024). Criticality therefore creates a condition in which technical knowledge about leverage can increase, rather than settle, the burden of political judgment.

Working Claim 1.2 (Dynamical leverage and normative warrant)**.** Evidence that a perturbation can strongly influence a near-critical trajectory provides information about leverage. It supplies only part of the justification required for intentional trajectory-shaping, because distributive effects, rights, legitimacy, reversibility, and uncertainty remain separate dimensions of judgment.

The distinction also changes the role of noise. A common policy intuition treats noise as a residual that should be filtered so that the underlying signal can be recovered. Nonlinear systems offer a wider range of possibilities. Weak noise can induce transitions between quasi-stable states, and stochastic resonance provides well-known examples in which a non-zero level of noise can improve the transmission or detection of weak signals under specific conditions (Astumian and Moss 1998; Forgoston and Moore 2018). These results do not license the political conclusion that noise is beneficial. They support a more limited methodological point: the function of a fluctuation cannot always be inferred from the fact that it is irregular, small, or difficult to predict.

Definition 1.3 (Perturbation–response relation)**.** For a perturbation η encountered in configuration (x, G), the *perturbation–response relation* denotes the trajectory-relevant change that emerges from the interaction among the perturbation, the receiving configuration, and subsequent feedback. Its political significance is therefore relational rather than a property of η considered in isolation.

The practical implication is a response repertoire broader than suppression. An irregular event can be observed, absorbed, buffered, redirected, temporarily decoupled, allowed to dissipate, selectively amplified, or met by strong intervention. The paper leaves the ordering of these responses open. Their relative desirability depends on the structure revealed by the perturbation, the possibility of cascading harm, the distribution of risk, the quality of available knowledge, and the ability to revise the response later. The phrase *epistem the noise* names this prior task of learning what the fluctuation is doing before its meaning is closed by an automatic threat classification.

This position resembles parts of the complexity-governance literature while placing greater emphasis on normative ambiguity. Work on governance and complex adaptive systems has already emphasized nonlinear change, thresholds, cascades, limited predictability, distributed institutions, learning, and adaptive capacity (Duit and Galaz 2008; May 2022). The present paper asks what follows when those properties are considered specifically near a sensitive transition. Adaptive capacity may protect a system, yet an adaptive intervention can also favor one trajectory over another. Polycentric response can create redundancy, yet it can also distribute responsibility in ways that make collective justification difficult. Early-warning information can support preparation, yet in reflexive political systems publication of that information can alter expectations and behavior. Near-critical response therefore combines dynamical and institutional questions with questions about authority and standing.

The resulting problem is especially visible in the relation between local knowledge and long-run consequence. A tangent-space approach to governance can be attractive under complexity because it limits action to locally supported information and encourages repeated updating. Near criticality, however, local validity and intervention safety can diverge. A small locally supported movement may place the system on another side of an attractor boundary, and the resulting trajectory may become difficult to reverse. The problem is not simply that the local estimate contains error. The system can amplify a small epistemic or normative mistake into a persistent historical change.

Working Claim 1.4 (Criticality boundary of local intervention)**.** Near a critical transition, the neighborhood within which local information is informative can be larger than the neighborhood within which intervention is readily reversible. Local epistemic support therefore need not provide a sufficient basis for local action.

This claim also changes the interpretation of better information. In ordinary control settings, improved estimates of sensitivity often increase the ability to act effectively. In political systems, improved identification of a high-leverage direction can instead increase responsibility for restraint. The actor learns that a small intervention may matter greatly; that knowledge also reveals that the consequences of error, exclusion, or unfair burden-sharing may be unusually large. Near-critical information can therefore support action, preparation, or deliberate non-forcing depending on the surrounding normative and institutional conditions.

The concept of *wuwei* provides one comparative resource for articulating this possibility. Interpretations of the *Laozi* emphasize that wuwei is not equivalent to total inactivity and can be read in terms of non-assertive or non-coercive action; water functions as a recurring image of yieldingness and strength (Chan 2025). The paper uses this material with two limits. First, Daoist political philosophy has its own historical and textual context, which should not be collapsed into a modern complexity framework. Second, water is a heuristic image rather than a model of political legitimacy. Its value here lies in challenging the assumption that response begins by removing every obstacle or fixing the final path in advance.

The water image helps distinguish force from configuration. A stone changes a flow without determining it independently. The flow responds to the stone, the terrain, velocity, prior channels, and subsequent changes in the environment. Political systems are far more reflexive and normatively structured, yet the image suggests a useful question: can governance sometimes improve its response by attending to the configuration through which a perturbation acquires force, rather than treating the perturbation as an isolated object to be eliminated? This orientation can be described provisionally through the language of *shi*, understood here only as a heuristic for directional propensity that emerges from configuration. The paper leaves a fuller historical and comparative treatment of this concept to later work.

The normative implications extend beyond the choice between intervention and restraint. Near criticality, a transition can change the set of futures that remain realistically reachable. Let $\Omega_i(t)$ denote, schematically, the future option space available to actor or group i at time t. An intervention can improve one present outcome while contracting $\Omega_i(t+1)$ for some actors. Justice may therefore concern not only the distribution of current benefits and burdens but also the distribution of future revisability and agency. This possibility gives reversibility a qualified normative importance. Reversible action is not always preferable, especially when urgent protection requires commitment, but the loss of reversibility increases the burden of justification under deep uncertainty.

Aesthetic judgment enters at a different stage of the analysis. Political actors sometimes perceive coherence, fragility, or emerging form before the consequences of a transition are fully manifest. Such perception can guide attention to possible trajectories and can help explain why some configurations appear salient before their effects are measurable. The paper nevertheless keeps aesthetic salience separate from moral endorsement. An emergent order can appear coherent or elegant while distributing burdens in ways that remain unjust. Aesthetics can therefore contribute to recognition of form without settling the ethics of trajectory selection.

These distinctions lead toward a response-centered orientation. Near criticality, political judgment concerns at least four layers. Dynamics concerns susceptibility, perturbation, coupling, feedback, and possible regime change. Epistemics concerns what can be known about those relations and how observation may itself become performative. Ethics and justice concern risk, burden, standing, agency, and the legitimacy of intentional trajectory-shaping. Aesthetics concerns the perception of emerging form before its normative meaning is settled. Treating these layers together reduces the temptation to translate a physical model directly into a political prescription.

Proposition 1.5 (Provisional response orientation)**.** *Where a political configuration is plausibly near a sensitive transition and knowledge remains incomplete, a defensible response framework may place weight on four capacities: characterization of perturbation–response relations, protection against severe and persistent harm, preservation of revisability where feasible, and explicit normative justification before intentional trajectory closure. This proposition is offered as an orientation for inquiry, not as a universal decision rule.*

The remainder of the paper will develop this orientation in stages. It will clarify the conceptual domain of near criticality, examine random noise as an epistemic and dynamical object, analyze response repertoires and the possibility of non-coercive participation, study hysteresis and trajectory closure, and then turn to justice, ethics, aesthetics, legitimacy, and institutional design. A later section will return to tangent-space governance and evidence-based policymaking to identify conditions under which locally valid evidence may be insufficient as a warrant for intervention. The final sections will state an empirical research agenda and preserve open problems rather than presenting a closed theory of near-critical politics.

Research Obligation 1.6 (Empirical and comparative validation)**.** Applications of this framework should specify the proposed transition mechanism, identify evidence for elevated susceptibility, distinguish endogenous from exogenous sources of fluctuation where possible, and state the limits of the analogy between the political case and the dynamical model invoked. Comparative work should also examine cases in which critical-transition language fails to improve explanation.

The paper's working position is therefore intentionally restrained. Criticality can identify conditions under which political response becomes unusually consequential, but it does not provide a political philosophy by itself. The central problem is how to participate in an evolving system when the future is partially open, perturbations can be amplified, intervention can become another perturbation, and the justice of alternative trajectories remains contested.

2 Conceptual Domain of Near Criticality

This section defines the conceptual domain in which the paper uses the language of near criticality. Its structural role is restrictive rather than expansive: it specifies the minimum dynamical features that make criticality a useful political vocabulary, distinguishes a near-critical region from an exactly identified critical point, and states conditions under which the language should be withheld. The discussion treats political criticality as a regime-sensitive hypothesis about response, feedback, and trajectory formation. It therefore avoids presuming that a political system possesses a single order parameter, a universal threshold, or a directly measurable analogue of a physical phase transition.

The starting intuition is comparatively modest. A political configuration can enter a condition in which the consequences of perturbation become unusually dependent on timing, coupling, and current structure. In such a condition, small differences in intervention or fluctuation can be associated with large differences in subsequent trajectories. Research on critical transitions in other complex systems provides formal examples of this general pattern, including changes in resilience, alternative stable states, and stochastic escape from locally stable regimes (Scheffer, Carpenter, et al. 2001; Scheffer, Bascompte, et al. 2009; Forgoston and Moore 2018). The present paper uses these results as conceptual resources. Political application requires an additional layer because political agents interpret signals, revise strategies, alter institutions, and react to attempts at steering. Complexity-oriented governance research accordingly treats institutions and governing actors as components of evolving systems rather than as external controllers (Duit and Galaz 2008; May 2022).

Definition 2.1 (Near-critical political region)**.** A *near-critical political region* is a provisional description of a configuration in which available evidence gives reason to consider the system's response to some class of perturbations unusually sensitive, trajectory dependent, and capable of producing a persistent reorganization of relations or institutions. The term identifies a region of elevated transition relevance. It does not require identification of an exact mathematical critical point.

Three features in this definition require emphasis. First, *sensitivity* concerns a relation between perturbation and response. A disturbance has no fixed political significance in isolation. Second, *trajectory dependence* concerns the possibility that nearby configurations can evolve along substantially different paths after apparently small differences in initial conditions, timing, or coupling. Third, *persistent reorganization* separates a candidate critical transition from an ordinary short-lived fluctuation. Persistence can arise through institutional reconfiguration, self-reinforcing expectations, network effects, path dependence, strategic adaptation, or other mechanisms that require independent empirical specification in each application.

A schematic state representation can make the distinction clearer. Let $x_t \in \mathcal{X}$ denote a high-dimensional political configuration and let η_t denote a perturbation drawn from a relevant class $\mathcal{E}$. The response over a finite horizon h can be represented as

$$\mathcal{R}_{t,h}(\eta_t) := d(\Phi_{t,t+h}(x_t; \eta_t), \Phi_{t,t+h}(x_t; 0)) \,, \tag{2}$$

where Φ denotes the evolving political dynamics and d is a problem-specific measure of configurational difference. Equation (2) is deliberately generic. It does not assume that a single Euclidean metric captures political distance. Its purpose is to make the object of interest explicit: near criticality concerns the response relation between a configuration and perturbations over a horizon, rather than the absolute magnitude of the perturbation alone.

This response-centered formulation also clarifies why the word *near* should remain relational. Proximity can refer to several distinct structures. A system may be close to a change in a slowly varying parameter, close in state space to a basin boundary, close to a network configuration that permits cascade propagation, or close in transition probability to a qualitatively new regime. These forms of proximity can diverge. A political configuration can appear macroscopically stable while its response structure

changes, and a visibly turbulent configuration can remain within the same broad institutional regime. Near criticality therefore cannot be inferred from surface volatility alone.

Remark 2.2 (Critical region and critical point)**.** The paper uses *critical region* when political evidence supports elevated sensitivity or transition relevance without supporting precise localization of a threshold. The stronger term *critical point* should be reserved for settings in which an explicit model and measurement procedure justify that precision. This distinction is especially important in reflexive social systems, where the governing relations themselves can evolve during observation and intervention.

The distinction has methodological consequences. An empirical study can make a credible case for near-criticality without claiming an exact threshold if it specifies the transition under consideration, identifies plausible feedback mechanisms, shows evidence of changing response or recovery properties, and documents the uncertainty surrounding alternative explanations. Conversely, calling a period "critical" solely because it appears historically dramatic adds little analytical value. Criticality language becomes useful when it changes what must be examined: perturbation–response relations, coupling, feedback, persistence, and the accessibility of alternative trajectories.

2.1 Susceptibility and finite political response

The first analytical component is susceptibility. In a smooth local model, susceptibility can be expressed through derivatives of response with respect to perturbation. Political applications require a wider interpretation because important responses can be discontinuous, delayed, or mediated by strategic reaction. Let u denote an intentional intervention and let Y denote a politically relevant observable. A local sensitivity quantity might be written as

$$\chi_{Y,u}(x_t) := \left. \frac{\partial \, \mathbb{E}[Y_{t+h} \mid x_t, u]}{\partial u} \right|_{u=0}, \tag{3}$$

when such a derivative is meaningful. Near a transition, however, the finite response represented in Equation (2) may be more informative than a derivative. A small perturbation can cross a boundary after which local linearization around the previous configuration has little value for the subsequent regime.

This distinction matters for the later critique of control-oriented reasoning. A large value of $\chi_{Y,u}$ indicates leverage with respect to a specified observable under a specified model. It does not establish the desirability of the resulting trajectory, the distribution of its consequences, or the safety of exploiting that leverage. Dynamical sensitivity is therefore an input to political judgment rather than a replacement for it.

Working Claim 2.3 (Susceptibility is relational)**.** The political significance of susceptibility should be indexed to a perturbation class, response variable, temporal horizon, and configurational context. A system should not be described as generically "highly susceptible" without specifying these dimensions.

This claim prevents an important category error. A configuration can be highly sensitive to financial expectations and comparatively insensitive to a legal reform; highly sensitive over weeks and comparatively stable over years; or sensitive in one region of a network while buffered in another. Near criticality can therefore be sectoral, relational, and scale dependent. The paper later uses this observation to motivate distributed sensing and response repertoires rather than a single system-wide policy instrument.

2.2 Alternative trajectories and attractor language

The second component is the possibility of qualitatively different persistent trajectories. The language of attractors is useful here when employed with care. An attractor in a formal dynamical system has a precise mathematical meaning. Political use is usually looser because institutions, preferences, networks,

and rules can change during the process being modeled. The paper therefore uses *attractor* as a structured hypothesis: a configuration may possess self-reinforcing mechanisms that make some patterns of interaction more likely to reproduce themselves once established.

Definition 2.4 (Political attractor hypothesis)**.** A *political attractor hypothesis* proposes that a family of political configurations contains feedback relations that increase the persistence or reproduction of that family after entry. The hypothesis requires specification of the relevant state description, feedback mechanisms, temporal scale, and conditions under which persistence weakens or fails.

This formulation preserves the central intuition without treating political history as motion on a fixed landscape. The landscape itself can change. A new institution can alter incentives; a public interpretation can change strategic expectations; an alliance can modify network structure; and an intervention can change the parameters that made the previous description useful. A political attractor is therefore better understood as a temporally situated pattern of self-reinforcement than as an immutable basin.

The same caution applies to hysteresis. In formal systems, hysteresis can mean that the path of return differs from the path of departure. Political applications may display analogous asymmetries when institutional dismantling, loss of trust, sunk investments, legal entrenchment, changed expectations, or new coalitions make restoration costly after a transition. The paper will use *hysteresis-like political persistence* where direct mathematical identification is unavailable. This wording keeps the empirical burden visible while retaining the normative implication: a transition can contract future revisability even when the initial intervention is small.

Proposition 2.5 (Transition relevance)**.** *Criticality language acquires political relevance when a perturbation can plausibly affect which self-reinforcing trajectory becomes accessible or persistent. Increased short-run volatility by itself provides insufficient support for this proposition.*

The proposition separates criticality from crisis rhetoric. Political crises can be severe while remaining within a single institutional regime. Conversely, a relatively quiet period can contain slowly changing couplings that increase the probability of a later transition. The analytic focus should therefore be placed on response structure and trajectory formation rather than on dramatic appearance.

2.3 Criticality as a configuration rather than a scalar condition

A third component concerns dimensionality. Political systems contain multiple actors, institutions, informational channels, material constraints, legal structures, and normative orders. Compressing this heterogeneity into one scalar "distance to criticality" can obscure the mechanism that matters. The paper therefore treats near criticality primarily as a property of a configuration and its perturbation–response relations.

Let G_t represent a network or relational structure, θ_t slowly changing institutional or material conditions, and H_t historically accumulated constraints. A generic representation is

$$x_{t+1} = F_t(x_t, G_t, \theta_t, H_t; \eta_t, u_t)\,. \tag{4}$$

The subscript on F_t is important. Political dynamics can change as actors learn, institutions are revised, and interpretations shift. Near criticality thus concerns a moving response field. Evidence obtained at one time can remain informative while losing intervention relevance after the governing relations change.

Equation (4) also introduces a distinction between endogenous fluctuation η_t and intentional intervention u_t for analytical convenience. The distinction is porous in practice. One actor's policy is another actor's perturbation; an intervention changes expectations; and a public warning about instability can alter the dynamics it describes. Later sections treat this reflexivity as a central limit of external-control metaphors.

Working Claim 2.6 (Configurational criticality)**.** For political analysis, evidence of near criticality should preferentially be expressed through identified configurations of coupling, feedback, and response rather than through an unqualified scalar label. Scalar indicators can remain useful when their relation to the underlying configuration is explicitly specified.

This claim also leaves room for partial criticality. A subsystem can approach a sensitive transition while other parts of a political order remain relatively stable. The resulting system can contain nested regions of susceptibility. Such heterogeneity matters normatively because intervention directed at one region can transmit effects through couplings that were not part of the initial policy target.

2.4 Noise, perturbation, and classification uncertainty

The fourth component concerns the status of noise. The term *noise* can refer to measurement error, unresolved micro-dynamics, exogenous stochastic forcing, endogenous fluctuation, or variation that appears irrelevant under a particular model. These meanings should remain distinct. The constructive role of stochastic fluctuation in physical and biological settings provides one reason to resist treating noise as equivalent to removable error (Astumian and Moss 1998; Forgoston and Moore 2018). Political analysis adds a further reason: what appears as noise under one explanatory frame can become politically meaningful after it couples to expectations, institutions, or networks.

Accordingly, the paper uses *perturbation* as the broader descriptive term and reserves *noise* for cases in which stochasticity, residual variation, or classificatory uncertainty is itself part of the problem. This terminological choice prevents the act of naming something "noise" from settling its normative status. A perturbation can be destructive, informative, generative, or largely irrelevant depending on the receiving configuration. Its significance emerges through relation.

This point also establishes the bridge to Section 3, where random noise will be treated as a dialectical object. Near criticality, the same fluctuation can obscure measurement while also probing susceptibility; it can destabilize an existing order while opening another; and attempts to remove it can alter the response capacity of the system. The conceptual domain of near criticality must therefore include uncertainty about the status of variation itself.

2.5 Working regime classification

For later analysis, it is useful to distinguish four working regimes without claiming sharp empirical boundaries. The classification is heuristic and should be revised when a case supplies better structure.

R1. Relatively buffered configuration. Relevant perturbations tend to produce bounded responses, and recovery or compensation remains readily available within the observed horizon.

R2. Elevated-susceptibility configuration. Response magnitude, duration, correlation, or propagation begins to increase for a specified perturbation class, while several future trajectories remain accessible.

R3. Transition region. Perturbations or interventions can plausibly affect entry into a different self-reinforcing configuration, and uncertainty concerning the consequences of small actions becomes politically salient.

R4. Post-transition configuration. New feedback relations have begun to reproduce a different regime, and restoration of the earlier configuration may require greater intervention, different conditions, or a new pathway.

The categories describe analytical attention rather than a deterministic sequence. A system can move between them, contain several simultaneously at different scales, or fail to exhibit a clear transition

despite increased volatility. Their purpose is to prevent the phrase *near criticality* from covering every period of uncertainty.

2.6 Conditions for withholding criticality language

A useful framework should specify its own domain of non-application. At least four situations call for withholding or weakening the criticality claim. First, observed changes may be adequately explained by approximately proportional responses within a stable regime. Second, the proposed alternative regimes may lack identifiable persistence mechanisms. Third, the relevant evidence may show disturbance without elevated susceptibility or trajectory bifurcation. Fourth, another vocabulary—for example bargaining, institutional change, organizational learning, or ordinary policy feedback—may explain the case with fewer unsupported assumptions.

Research Obligation 2.7 (Domain discipline)**.** An empirical application should state what would count against the near-critical interpretation. At minimum, it should identify plausible non-critical explanations, specify the perturbation class and response horizon, and describe the evidence that would support a return to ordinary regime-specific analysis.

This obligation protects the framework from retrospective inflation. Once a major political transformation has occurred, it is easy to redescribe every preceding fluctuation as a warning signal. A disciplined analysis should preserve the possibility that the transition was not detectable in advance, that the criticality analogy is weak, or that the observed sequence arose from a different mechanism. Historical importance does not retroactively establish near-critical observability.

The conceptual domain established in this section is therefore intentionally bounded. Near criticality refers to a configuration in which perturbation, response, feedback, and trajectory selection acquire unusual significance. The concept does not supply a universal metric of political instability, and it does not by itself authorize intervention. Its value lies in changing the questions posed to a political system: which relations are becoming sensitive, which fluctuations reveal or alter those relations, which trajectories can become self-reinforcing, and which future options may contract after a transition. Section 3 develops the first major consequence of this reframing by examining random noise as an epistemic, dynamical, and political object.

3 Random Noise and Political Significance

This section develops the paper's treatment of random noise as an epistemic, dynamical, and political object. Its structural role is to separate three questions that are often compressed into one: what variation is being observed, what that variation does when it enters a particular configuration, and what kind of response, if any, is warranted. The discussion proceeds from stochastic dynamics to political interpretation without assuming a direct identity between the two domains. The central methodological commitment is that the initial magnitude or apparent irregularity of a fluctuation supplies only partial information about its political significance. Near criticality, significance depends increasingly on coupling, susceptibility, persistence, and the response capacity of the receiving configuration.

Figure 2 summarizes the shift introduced here. A fluctuation is treated neither as self-interpreting evidence nor as an intrinsically harmful deviation. Its significance is mediated by a response field. The same perturbation class may dissipate in one configuration, reveal latent susceptibility in another, and contribute to a transition in a third. This response-centered view also prepares the argument of Section 4, where the perturbation–response pair becomes the primary analytical unit.

3.1 Analytical statuses of noise

The term *noise* carries several analytically distinct meanings. In an empirical measurement, noise can designate error introduced by an instrument or sampling procedure. In a model, it can designate unresolved variation left outside the state description. In a stochastic dynamical system, it can represent external forcing or internally generated randomness. In political interpretation, the same word is also used rhetorically for events judged irrelevant to a preferred explanatory frame. These uses should be separated before any policy significance is assigned.

A simple decomposition,

$$y_t = s_t + \epsilon_t, \tag{5}$$

can be useful when the distinction between signal and error is supplied by a well-specified measurement problem. Political systems often make the decomposition more difficult. A residual term can contain measurement error, unobserved strategic interaction, local experimentation, exogenous shock, endogenous fluctuation, or an emerging pattern that the present model has not yet represented. Treating the residual as a homogeneous category can therefore combine ontological uncertainty with epistemic limitation.

Definition 3.1 (Analytical noise)**.** For the purposes of this paper, *analytical noise* denotes variation whose causal status, informational relevance, or dynamical role is unresolved within the current description. This definition leaves open whether the variation is stochastic in the world, unmeasured structure, model residual, or a mixture of these sources.

The definition is intentionally modest. It avoids converting an epistemic classification into an ontological conclusion. An analyst may have good reason to model a disturbance stochastically while still recognizing that later inquiry could reveal structure at another scale. Conversely, apparent patterns can arise through stochastic processes and need not acquire political meaning merely because they are narratively interpretable. Near-critical analysis therefore requires uncertainty about the classification of variation itself.

This distinction also clarifies the phrase used throughout the paper, *epistemics of noise*. The phrase refers to practices for identifying the source, correlation structure, scale, persistence, and response relevance of a fluctuation before assigning it a fixed political role. Such practices can include statistical characterization, network observation, qualitative process tracing, comparison across scales, and explicit representation of alternative explanations. The methodological point is plural: near-critical fluctuations can require more than one observational language.

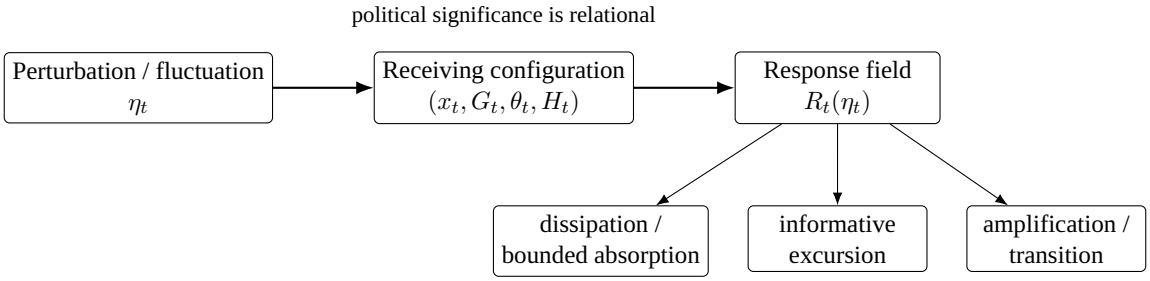

Figure 2: A response-centered interpretation of perturbation. The same class of fluctuation can have different consequences when it enters different configurations. The figure is conceptual and does not imply that the three response modes are exhaustive or mutually exclusive.

3.2 Constructive stochasticity and domain restraint

The physical literature supplies a useful corrective to the intuition that randomness always degrades order. Work on stochastic resonance shows that, under specific nonlinear conditions, noise can assist the transmission or amplification of weak signals (Gammaitoni et al. 1998). Reviews of fluctuation-driven transport likewise describe circumstances in which random motion contributes to directed behavior rather than simply obscuring it (Astumian and Moss 1998). In stochastic dynamical systems, even weak noise can also induce transitions between quasi-stable states (Forgoston and Moore 2018). These results establish a limited but important lesson: the dynamical role of noise depends on system structure.

The political use of this lesson requires restraint. The paper does not infer that political "noise" performs stochastic resonance in the technical sense, or that social variation should be celebrated because constructive stochastic phenomena exist in physics. The transfer is conceptual and conditional. It supports a narrower proposition: random variation can interact with nonlinear structure in more than one way, so a political theory that equates fluctuation with removable error can miss relevant response mechanisms.

Proposition 3.2 (Non-equivalence of noise and harm)**.** *The classification of a fluctuation as noise does not by itself establish that reducing the fluctuation improves the political system. A separate account of its dynamical role, distributional effects, and relation to system viability is required.*

This proposition matters especially near criticality because small variations can acquire system-level relevance through amplification. A regulator who suppresses every observable deviation can reduce short-run variance while also removing exploratory variation, concealing susceptibility, or changing the feedback structure that previously dissipated disturbance. These possibilities are hypotheses to be investigated rather than assumptions in favor of permissiveness. The symmetrical error is to romanticize disturbance and treat every disruption as generative. Constructive and destructive roles remain empirical and normative questions.

The distinction can be represented by a response functional

$$\mathcal{R}_h(\eta_t \mid \mathcal{C}_t) := \mathbb{P}(X_{t+h} \in \cdot \mid \eta_t, \mathcal{C}_t)\,, \tag{6}$$

where $\mathcal{C}_t$ denotes the relevant political configuration. Equation (6) emphasizes that the object of interest is conditional response over a horizon. The same distribution of perturbations can generate different response distributions as $\mathcal{C}_t$ changes.

3.3 Fluctuation as an epistemic probe

Near critical regions, fluctuations can carry information about the system's response structure. The early-warning literature on critical transitions uses this idea in a formal way. In classes of systems displaying

critical slowing down, recovery from perturbation becomes slower as a threshold is approached, and changes in variance or autocorrelation can sometimes provide warning signals (Scheffer, Bascompte, et al. 2009). These indicators are neither universal nor sufficient for political diagnosis, but the underlying epistemic idea is valuable: a perturbation can reveal properties that are difficult to infer from a static state description.

Political analysis can generalize this idea without assuming identical mathematics. Consider two configurations with similar aggregate indicators. A modest shock produces rapid local compensation in the first configuration and prolonged cross-sector propagation in the second. The disturbance has functioned as a probe of coupling and recovery. The politically relevant information lies partly in the response, not solely in the pre-shock state or the shock magnitude.

Working Claim 3.3 (Response-based evidence)**.** Near criticality, evidence about how a system recovers from, propagates, or reorganizes around perturbations can be more informative about transition risk than evidence about the perturbation magnitude considered in isolation.

The claim motivates a different observational posture. Analysts should record response duration, propagation paths, changes in correlation, institutional adaptation, expectation shifts, and the emergence of new self-reinforcing relations. A fluctuation that appears trivial under an event-centered dataset can become important once its response trajectory is observed. Conversely, a large event that is rapidly absorbed may provide evidence of buffering capacity rather than imminent transition.

This perspective also complicates intervention. If naturally occurring fluctuations provide information about resilience and susceptibility, immediate suppression can remove an observational opportunity. Delayed response can sometimes increase learning, while delay can also permit amplification. The choice therefore depends on reversibility, plausible harm, and the expected value of additional information. Sections 6 and 13 return to this tension from epistemic and ethical perspectives.

3.4 Noise, organization, and recursive classification

The paper uses the term *dialectical object* in a restricted sense to capture a recursive relation. Noise can disturb an existing organization; the response to disturbance can generate another organization; and the new organization can subsequently redefine which variations count as relevant, ordinary, or noisy. The categories of signal and noise can therefore change as the system changes.

A stylized representation is

$$\mathcal{C}_t, \eta_t \longmapsto \mathcal{C}_{t+1} \longmapsto \Pi_{t+1}(\eta), \tag{7}$$

where Π_{t+1} is a later classificatory scheme for variation. The equation is not intended as a deterministic law. It marks the possibility that a political order helps determine which deviations become visible or significant, while deviations can participate in transforming that order.

This recursive structure matters for institutional design. A governance system that treats its current classification scheme as permanent may suppress variation that would become meaningful under an emerging configuration. A system that constantly revises its categories can generate instability of a different kind. Near-critical response therefore requires a balance between classification continuity and the capacity to recognize novelty.

Remark 3.4 (Historical dependence)**.** The political meaning of noise can be historically dependent. A type of variation that once dissipated locally can acquire greater significance after changes in technology, network structure, institutional coupling, or collective expectation. Historical recurrence of an event therefore does not guarantee recurrence of its system-level effect.

This remark connects random noise to the paper's larger concern with historical becoming. Political actors can learn from past responses, but the learning process itself changes the response field. Near-critical analysis should therefore avoid a simple archive of "shock sizes" detached from the evolving configuration in which those shocks occurred.

3.5 Response before suppression

The preceding analysis supports a response-centered orientation. When a fluctuation appears, suppression is one possible response among several. Other modes include observation, buffering, dissipation, decoupling, redirection, selective accommodation, and, in carefully justified cases, amplification. The selection among these modes requires information about the receiving field and normative judgment about the consequences of alternative trajectories.

This orientation can be summarized by distinguishing perturbation magnitude from response significance. Let

$$\Sigma_t(\eta) := \Phi(\|\eta\|, \chi_t, G_t, \tau_t, \rho_t, \mathcal{V}_t)\,, \tag{8}$$

where χ_t denotes relevant susceptibility, G_t coupling structure, τ_t response persistence, ρ_t reversibility or return conditions, and $\mathcal{V}_t$ a provisional representation of viability and normative stakes. Equation (8) is schematic. Its purpose is to prevent $\|\eta\|$ from serving as a proxy for significance. A small disturbance can become politically consequential in a highly susceptible and tightly coupled configuration, while a larger one can remain bounded when buffers and alternative pathways are available.

Working Claim 3.5 (Response priority)**.** Near criticality, the first analytical task is to characterize the relation between fluctuation and response. Intervention choice follows from that characterization together with ethical, distributive, and legitimacy considerations.

The phrase *response priority* does not prescribe passivity. Some perturbations require rapid containment because delay could create severe or irreversible harm. The claim concerns analytical ordering: the desirability of suppression cannot be read directly from the existence of fluctuation. Even an urgent response benefits from distinguishing the disturbance, its propagation mechanism, and the channels through which intervention can itself become a new perturbation.

3.6 Implications for near-critical political inquiry

Random noise therefore enters the theory in four connected capacities. First, it is an *epistemic challenge*: analysts may be uncertain whether observed variation is error, stochastic forcing, unresolved structure, or an emerging pattern. Second, it is a *dynamical input*: stochastic fluctuations can be amplified, dissipated, or involved in state transitions. Third, it is an *epistemic resource*: system responses to fluctuation can reveal susceptibility, coupling, and recovery properties. Fourth, it is a *political object*: institutions must decide how much variation to absorb, permit, redirect, or constrain, and those decisions distribute risk and future possibility unevenly.

These capacities explain why random noise is dialectically important to the paper. Order and noise cannot be assigned stable political value in advance. An order can be sustained by suppressing fluctuation, by metabolizing it, or by allowing limited reorganization; a fluctuation can damage an order, expose its fragility, or participate in another form of organization. The relevant judgment concerns the relation among disturbance, response, trajectory, and the people who bear their consequences.

Research Obligation 3.6 (Noise characterization)**.** Applications of the framework should document how a fluctuation is classified, what alternative classifications remain plausible, which response variables are observed, and which normative assumptions enter the decision to suppress, permit, or otherwise respond to it.

This obligation provides a bridge from descriptive dynamics to political responsibility. Calling variation "noise" should not settle the policy question. Near criticality, the same label can conceal measurement uncertainty, unmodeled relation, transition-inducing stochasticity, or adaptive variation. A responsible analysis keeps these possibilities visible long enough to examine how the system actually responds. Section 4 therefore shifts the unit of analysis from the fluctuation alone to the perturbation–response relation and develops a vocabulary for response modes, propagation, and trajectory effects.

4 Perturbation–Response Relations

This section develops the perturbation–response relation as the paper's primary unit of dynamical analysis. Its structural role is to move from the classification of random noise in Section 3 toward a framework for describing what happens after a disturbance enters a political configuration. The discussion has four objectives. First, it defines a response field that represents the dependence of consequences on perturbation, receiving configuration, and temporal horizon. Second, it separates response magnitude from response form by distinguishing absorption, amplification, propagation, redirection, and transformation. Third, it introduces memory, heterogeneity, and network structure as reasons why apparently similar disturbances can produce different political trajectories. Fourth, it clarifies the limited role of response information in political judgment: evidence about sensitivity or leverage describes a dynamical relation, while decisions about intervention continue to require independent normative justification.

The general language of response has a long history in statistical physics and nonequilibrium theory. Linear response theory represents the effect of a small external perturbation through response functions and susceptibilities, while also relating response to fluctuations under specified assumptions (Kubo 1982). Network research extends the same intuition by showing that collective response can depend strongly on topology, the location of a stimulus, and whether a perturbation affects units or their interactions (Bar-Yam and Epstein 2004; Manik et al. 2017). The present paper does not transfer those formalisms directly into politics. It uses them to discipline a more general proposition: the political significance of a disturbance depends on how a historically situated and relationally structured system receives it.

Definition 4.1 (Political response field)**.** For a political configuration described by state variables x_t, relational structure G_t, and relevant history H_t, the *political response field* is a provisional map from a class of perturbations to distributions over trajectory-relevant changes across a specified horizon. It can be written schematically as

$$\mathfrak{R}_{t,h} : \eta \longmapsto \mathcal{P}(\Delta\Gamma_{t:t+h} \mid x_t, G_t, H_t, \eta)\,, \tag{9}$$

where $\Delta\Gamma_{t:t+h}$ denotes a change in the subsequent trajectory and $\mathcal{P}$ represents uncertainty over that change.

Equation (9) makes three restrictions explicit. Response is indexed to a perturbation class, because the same system can be sensitive to one kind of disturbance and buffered against another. Response is indexed to a horizon, because a disturbance can generate a visible short-run displacement and little persistent change, or a small immediate displacement followed by slow amplification. Response is also conditioned on history. Political configurations carry institutions, expectations, remembered conflicts, commitments, network relations, and accumulated asymmetries that shape what a new disturbance can become.

Figure 3 summarizes the section's central shift. The perturbation is treated as an input to a response relation rather than as a self-interpreting political event. A disturbance described as minor at the point of entry can propagate through highly susceptible relations; a large shock can remain bounded when redundancies, buffers, or weak couplings limit transmission. Research on networked systems provides formal examples in which response depends on topology and stimulus location rather than on stimulus size alone (Bar-Yam and Epstein 2004; Manik et al. 2017). Political application adds reflexive actors and contested meanings, but the analytical caution remains useful: event magnitude cannot substitute for response analysis.

4.1 Directional susceptibility and finite response

A response field contains direction as well as magnitude. Let $Y = (Y_1, \ldots, Y_m)$ denote a set of observables and $\eta = (\eta_1, \ldots, \eta_k)$ a perturbation vector. When a local derivative is meaningful, a directional

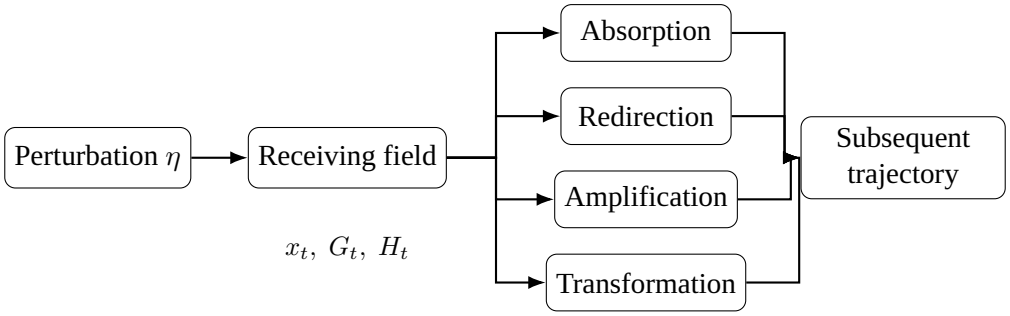

Figure 3: A perturbation acquires trajectory significance through the receiving field. The same disturbance can be absorbed, redirected, amplified, or involved in transformation depending on state, coupling, history, and subsequent feedback. The categories are analytical modes and can overlap in empirical cases.

susceptibility matrix can be represented as

$$\boldsymbol{\chi}_{t,h} = \left[\frac{\partial\, \mathbb{E}[Y_i(t+h)]}{\partial \eta_j} \right]_{i,j}. \tag{10}$$

A large entry identifies a direction in which a small perturbation is associated with a comparatively large expected change in a selected observable. Network susceptibility research makes a related distinction between responses to changes in node properties and responses to changes in interactions (Manik et al. 2017). For political systems, the distinction suggests that sensitivity can reside in actors, relations, institutional rules, or the connections through which information and resources travel.

The matrix in Equation (10) remains a local object. Near a transition, finite perturbations can carry a system outside the region in which the derivative describes later dynamics. The paper therefore uses a finite amplification ratio as a complementary descriptive quantity:

$$\mathcal{A}_{t,h}(\eta) := \frac{d(\Phi_{t,t+h}(x_t;\eta), \Phi_{t,t+h}(x_t;0))}{\|\eta\|}, \tag{11}$$

where d denotes a problem-specific measure of configurational difference. The ratio is schematic and can lose meaning when the numerator and denominator belong to incomparable spaces. Its purpose is conceptual: elevated response can appear as a disproportion between initiating disturbance and later divergence, while the content of that divergence still requires substantive interpretation.

Working Claim 4.2 (Directional response)**.** Near-critical sensitivity should be described directionally. Evidence that a system is responsive to one perturbation channel does not establish comparable responsiveness to other channels, and an aggregate sensitivity measure can hide important asymmetries among actors, relations, and scales.

The claim matters politically because a system can combine fragility and robustness. A network can be sensitive to perturbations entering selected nodes or relations and remain resistant to many other disturbances. Bar-Yam and Epstein formalize this coexistence in a network model in which sensitivity to selected stimuli and robustness to others can both contribute to adaptive function (Bar-Yam and Epstein 2004). Political systems can display analogous heterogeneity without sharing the same mechanism. A diplomatic channel, market infrastructure, media network, coalition structure, or constitutional rule can create a high-susceptibility pathway inside an otherwise stable configuration.

4.2 Response modes and trajectory effects

Response magnitude alone remains incomplete because large changes can take different forms. The framework therefore distinguishes five analytical modes. They describe trajectory effects and can coexist

in the same episode.

Definition 4.3 (Response modes)**.** Given a perturbation η and receiving configuration $\mathcal{C}_t$, the following modes describe different relations between disturbance and subsequent dynamics:

(i) *absorption*: the perturbation produces a bounded displacement and the configuration returns toward a comparable region of activity;

(ii) *amplification*: subsequent feedback enlarges the disturbance's trajectory effect;

(iii) *propagation*: effects travel across actors, sectors, networks, or scales beyond the initial point of entry;

(iv) *redirection*: the disturbance is converted into another form of response, shifting the channel, location, or domain in which its effects appear;

(v) *transformation*: the response alters the relations or rules that shape later responses, thereby modifying the response field itself.

Absorption is related to recovery, although recovery need not imply exact return to a previous state. Work on critical slowing down shows that recovery from small perturbations can become slower as some systems approach certain classes of transition (Nes and Scheffer 2007; Scheffer, Bascompte, et al. 2009). Political recovery requires a broader description because institutions and expectations can be changed by the disturbance even when aggregate indicators return toward their previous values. A response can therefore appear stable at one level while leaving a persistent relational trace at another.

Amplification concerns positive feedback. An initial disturbance can alter expectations, behavior, resource allocation, or network connections in ways that increase subsequent change. Propagation concerns spatial, organizational, or relational transmission. The two often interact: a perturbation can spread without growing at each step, or it can amplify locally without spreading widely. Redirection captures another possibility. A system can transform a political disturbance into legal contestation, market adjustment, institutional reform, public narrative, or another channel whose dynamics differ from the initial event. Transformation is the strongest mode because it changes the rules of later reception. After a constitutional reform, alliance reconfiguration, technological shift, or institutional collapse, the same future perturbation can enter a different response field.

Remark 4.4 (Response modes and value)**.** The response modes carry no fixed normative ordering. Absorption can protect a valuable institution or preserve an unjust lock-in. Amplification can generate destructive cascades or accelerate a broadly supported transition. Redirection can diffuse conflict or displace burdens onto less visible actors. Transformation can expand or contract future agency. Normative evaluation therefore enters after the response relation has been described, while also shaping which response variables receive attention.

This remark is central to the paper's separation of dynamics and judgment. A response vocabulary helps identify what a system does with disturbance. It does not supply the ranking of those outcomes. The distinction prevents the concept of "adaptive response" from acquiring an automatic positive valence. Adaptation by one actor can create vulnerability for another; institutional survival can maintain unequal distributions; and systemic flexibility can be purchased through concentrated exposure elsewhere.

4.3 Memory and delayed response

Political response is historically conditioned. The same perturbation applied to similar observable states can produce different outcomes when the underlying histories differ. A minimal representation extends the state variable with a history-dependent term:

$$\mathfrak{R}_{t,h} = \mathfrak{R}_h(x_t, G_t, H_t), \tag{12}$$

where H_t can include accumulated commitments, remembered harms, policy precedents, institutional learning, expectations, and prior interventions. This representation allows response to display memory without presuming a specific mathematical form of hysteresis. Section 5 will develop the stronger case in which entry and return follow different conditions.

Memory also creates delayed response. A disturbance can initially appear absorbed while modifying variables whose consequences emerge later. Conversely, an immediate spike can dissipate without reorganizing the longer trajectory. Near-critical analysis should therefore resist equating short-run volatility with long-run transition relevance. Horizon selection becomes part of the research design, and multiple horizons can be necessary when political effects move through institutions at different speeds.

Proposition 4.5 (Temporal non-equivalence)**.** *Two perturbations that produce similar immediate responses can remain trajectory-distinct when they create different changes in expectations, relations, institutional memory, or future susceptibility. Short-horizon response equivalence therefore does not establish long-horizon trajectory equivalence.*

The proposition motivates later attention to reversibility. A policy can restore a visible variable while leaving behind a changed distribution of trust, authority, debt, legal precedent, or strategic expectation. A return in observables can coexist with a failure to return in the response field. That possibility is one reason near-critical political analysis requires more than monitoring the state vector alone.

4.4 Heterogeneous response and cross-scale propagation

Political systems consist of interacting actors and institutions with unequal positions, capacities, and exposures. An aggregate response can therefore hide local amplification. Let $r_i(\eta)$ denote a response associated with actor or subsystem i. A small average response,

$$\bar{r}(\eta) = \frac{1}{n}\sum_{i=1}^{n} r_i(\eta), \tag{13}$$

can coexist with large responses in a subset of units. The distribution of r_i and the network through which these responses interact can matter more than the mean. Network studies formalize analogous dependence of collective response on connectivity and perturbation location (Bar-Yam and Epstein 2004; Manik et al. 2017).

For politics, heterogeneous response has at least three implications. First, local vulnerability can remain invisible in aggregate indicators until propagation begins. Second, actors can strategically redirect perturbations, passing costs or risks through institutional and economic relations. Third, heterogeneity can provide buffering when different units respond through partially independent mechanisms. The same diversity that complicates prediction can reduce correlated failure under some configurations. These possibilities motivate a later treatment of response repertoires and institutional architecture rather than a single system-wide control variable.

Cross-scale propagation adds another difficulty. A local event can remain local, move horizontally through a network, or alter higher-level rules that feed back onto local actors. Conversely, a system-level intervention can create local responses that aggregate into an unintended macro trajectory. The response field is therefore potentially multilevel:

$$\eta^{(\ell)} \longmapsto \left\{\Delta\Gamma^{(\ell-1)}, \Delta\Gamma^{(\ell)}, \Delta\Gamma^{(\ell+1)}\right\}, \tag{14}$$

where ℓ denotes a provisional level of organization. The notation is illustrative. Political scales can overlap, and actors can participate in several relational structures simultaneously.

Working Claim 4.6 (Response-field heterogeneity)**.** Near-critical political analysis should examine the distribution and coupling of responses rather than relying exclusively on aggregate indicators. A sys-

tem can appear stable in the mean while containing localized high-susceptibility pathways capable of transmitting or amplifying disturbance.

4.5 Response information and political action

The response field provides information about how disturbance can matter. It can reveal susceptible directions, slow recovery, potential propagation channels, and transformations in the system's own capacity to receive later perturbations. Such information can improve diagnosis. It can also create a powerful temptation to identify high-susceptibility directions with desirable points of intervention.

The paper separates these steps. Suppose analysis indicates that a small intervention u produces a large response in direction v. The finding can be represented schematically as

$$\|\mathfrak{R}_{t,h}(u)\| \gg \|u\|. \tag{15}$$

Equation (15) identifies leverage. It leaves open whether the induced trajectory is just, whether affected actors have standing in the decision, whether risks are fairly distributed, whether the response is reversible, and whether publication or use of the information changes the response field. Near criticality, response information can therefore increase both capability and responsibility.

Working Claim 4.7 (Response information and restraint)**.** Evidence of high susceptibility can support preparation, buffering, selective intervention, or restraint. The direction of action cannot be inferred from susceptibility alone because the same information that identifies leverage also identifies the possibility of disproportionate and persistent consequence.

This claim provides the bridge to the paper's later political-philosophical analysis. A response-centered framework avoids an automatic movement from observation to optimization. It allows a perturbation to be examined as a probe of the receiving configuration, an event with distributive consequences, and a possible participant in trajectory formation. It also keeps open the possibility that an appropriate political response consists in buffering one channel while leaving another undisturbed, delaying a decision until additional structure becomes visible, or declining to exploit a high-leverage direction.

Research Obligation 4.8 (Response specification)**.** Empirical applications should specify the perturbation class, receiving configuration, temporal horizon, response variables, propagation pathways, recovery criteria, and distribution of response across affected actors. Claims of high sensitivity should distinguish local derivatives from finite trajectory changes and should document uncertainty about alternative response mechanisms.

The perturbation–response relation therefore provides a bridge between noise and trajectory. Section 3 established that irregular variation can be epistemically ambiguous and dynamically consequential. The present section shows that its significance is generated through response. Section 5 turns to the next problem: some response patterns become self-reinforcing, entry into a new regime can alter the conditions of return, and an intervention can participate in closing a wider field of possible futures into a more persistent trajectory.

5 Attractors, Hysteresis, and Trajectory Closure

This section develops the temporal structure that makes near-critical response politically consequential. Its role is to connect the perturbation–response relations of Section 4 to persistence after a transition. The discussion has four objectives. First, it refines the paper's use of attractor language by identifying self-reinforcing mechanisms rather than treating political history as motion on a fixed landscape. Second, it introduces hysteresis as a model of return asymmetry: conditions sufficient to enter a regime can differ from those required to leave it. Third, it defines *trajectory closure* as a contraction in the set of practically reachable future configurations after commitments, feedback, and accumulated costs begin to reinforce a path. Fourth, it separates dynamical persistence from normative value. A regime can be persistent without being just, and an easily reversible configuration can still be normatively objectionable.

The physical and ecological literature supplies precise examples of alternative stable states, basin crossing, and hysteresis. In several ecological models, gradual changes in parameters can produce abrupt transitions to contrasting states, while reversal of the original parameter change need not immediately restore the previous state (Scheffer, Carpenter, et al. 2001; Beisner et al. 2003). Research on recovery from perturbations similarly relates reduced recovery rates to reduced resilience near some catastrophic shifts (Nes and Scheffer 2007). These results provide conceptual discipline rather than direct political evidence. Political institutions, expectations, network relations, and normative commitments can change while a transition unfolds. For this reason, the present paper treats attractors and hysteresis as mechanism-sensitive hypotheses whose empirical content must be specified in each application.

Figure 4 separates two ideas that are often merged. Hysteresis concerns asymmetry between entry and return. Trajectory closure concerns the changing accessibility of alternative futures after a path begins to reproduce itself. They can occur together, but neither implies the other in all systems. A political configuration can display strong path dependence without possessing a clean bifurcation diagram, and a formal dynamical model can display hysteresis without containing the institutional or normative features that make political lock-in significant.

5.1 Self-reinforcement and attractor hypotheses

The paper's attractor vocabulary begins from persistence rather than geometry. Suppose a family of configurations $\mathcal{A}$ contains mechanisms that make continued residence in that family more likely after entry. A provisional persistence relation can be represented as

$$\Pr(x_{t+h} \in \mathcal{A} \mid x_t \in \mathcal{A}, H_t) > \Pr(x_{t+h} \in \mathcal{A} \mid x_t \notin \mathcal{A}, H_t)\,, \tag{16}$$

for an explicitly defined horizon and history H_t. Equation (16) does not define a mathematical attractor. It identifies an empirical pattern for which an attractor hypothesis may be worth testing. The analytical burden then lies in identifying mechanisms that account for persistence.

Political self-reinforcement can arise through several mechanisms. Institutional investment can increase switching costs. Coordination can make an established rule more attractive because others already rely on it. Organizations can acquire resources that enable them to defend the arrangements from which those resources arose. Expectations can become mutually confirming when actors adjust strategies to what they believe others will continue to do. Legal or administrative infrastructures can make alternative procedures costly to reconstruct. Historical-institutionalist work on increasing returns and path dependence emphasizes precisely the importance of timing, sequence, feedback, and rising reversal costs in political development (Pierson 2000). The present framework places those mechanisms alongside near-critical sensitivity: a perturbation can matter greatly when it occurs before a self-reinforcing configuration becomes expensive to leave.

Definition 5.1 (Self-reinforcing political trajectory)**.** A *self-reinforcing political trajectory* is a temporally situated path in which consequences of prior movement alter incentives, expectations, capabilities,

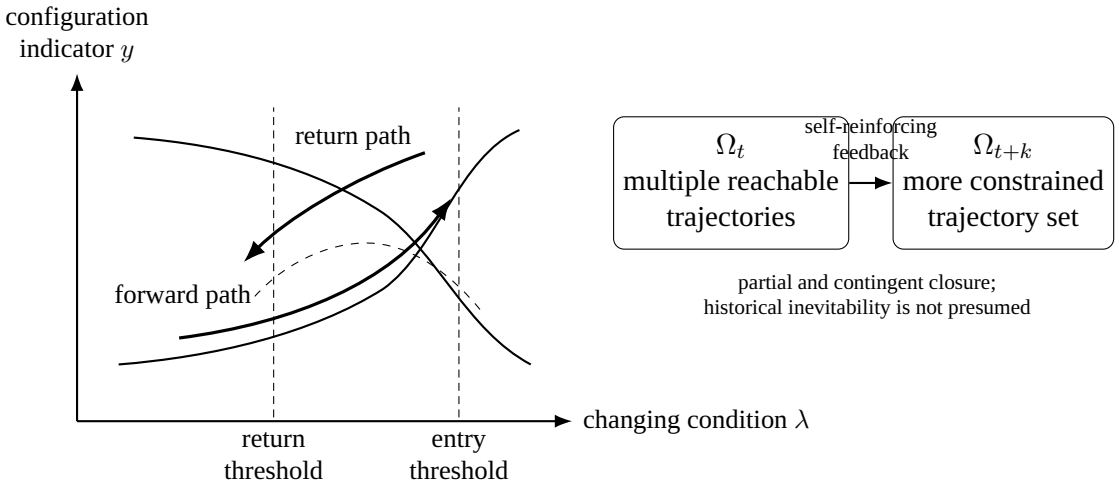

Figure 4: Two related meanings of persistence. The left panel gives a schematic hysteresis relation in which the condition associated with entry differs from the condition associated with return. The right panel represents trajectory closure as a possible contraction of practically reachable futures after feedback and accumulated commitments. The figure is conceptual; political applications require independently specified mechanisms and measures.

institutions, or relations in ways that increase the relative persistence of continued movement within the same broad configurational family. The definition leaves open whether reinforcement is efficient, legitimate, just, or desirable.

This definition avoids equating persistence with equilibrium. A trajectory can remain internally dynamic while reproducing a broader configuration. Actors can change strategies, institutions can be revised, and conflicts can continue even when those changes reinforce the same larger pattern. The relevant question is whether the evolving configuration changes the costs and probabilities of future movement in a path-dependent direction.

Remark 5.2 (Attractors and endogenous landscapes)**.** The political landscape should not be treated as fixed while actors move across it. Entry into a trajectory can alter the landscape itself by changing rules, resources, expectations, networks, and available categories of action. An attractor hypothesis in politics therefore concerns co-evolution between state and landscape rather than motion within an invariant potential surface.

This remark has consequences for near-critical response. An intervention that appears small in the pre-transition configuration can participate in creating a new response field. Subsequent interventions are then evaluated within a changed set of relations. The political significance of a transition therefore includes the possibility that the transition changes the mechanisms through which later perturbations are received.

5.2 Hysteresis and return asymmetry

Hysteresis provides a precise way to distinguish entry conditions from return conditions. In a stylized system controlled by a parameter λ, a forward transition can occur at λ_{in}, while return to the previous state occurs only after the parameter is moved past a different value λ_{return}. The difference

$$\Delta_{\text{hys}} := |\lambda_{\text{in}} - \lambda_{\text{return}}| \tag{17}$$

can represent the width of a hysteresis region in a model where such a scalar parameter is meaningful. Ecological work on alternative stable states uses this form of return asymmetry to distinguish a

reversible displacement from a regime shift in which restoring the earlier driver value can be insufficient for restoration (Scheffer, Carpenter, et al. 2001; Beisner et al. 2003).

Political applications should use a broader concept because a single control parameter will often be unavailable. The relevant asymmetry can instead be expressed in terms of conditions or costs. Let $C_{A\to B}$ denote the resources, institutional changes, or coordinated actions associated with moving from a configurational family A toward B, and let $C_{B\to A}$ denote the corresponding requirements for return. A transition can display politically important return asymmetry when

$$C_{B\to A} \not\approx C_{A\to B}, \tag{18}$$

with the direction and magnitude established empirically rather than assumed. The asymmetry can arise because the transition modifies institutions, distributes new resources, creates expectations, dissolves organizations, or changes the credibility of previous commitments.

Definition 5.3 (Political return asymmetry)**.** *Political return asymmetry* denotes a condition in which reversing the initial intervention or restoring an earlier observable condition is insufficient to reproduce the earlier political configuration, because the transition has altered relevant relations, institutions, expectations, capabilities, or accumulated commitments.

This definition is intentionally wider than formal hysteresis. It captures a political analogue without claiming that every instance of path dependence is a bifurcation. The distinction also protects the analysis from a common error: observing persistence after a shock does not by itself establish hysteresis. Persistent change can arise through exogenous conditions that remain in place, through slow adjustment, or through mechanisms unrelated to alternative stable states. Empirical work should therefore identify the mechanism that produces return asymmetry.

Working Claim 5.4 (Asymmetry of leverage and reversal)**.** Near a transition, a comparatively small intervention can sometimes have large trajectory consequences while the resources required to reverse the resulting configuration later become substantially larger. Evidence of low entry cost therefore provides little information about the cost or feasibility of return.

The claim is conditional. Its importance lies in the mismatch between two quantities that a control-oriented analysis can easily conflate. The first is the marginal effort required to influence a sensitive system at time t. The second is the effort required to reopen alternatives after self-reinforcement has altered the response field. Near criticality, low leverage cost can coexist with high reversal cost. This possibility will later ground the paper's analysis of responsibility and irreversibility.

5.3 Trajectory closure and practical reachability

Return asymmetry concerns movement back toward a previous configuration. Trajectory closure concerns something wider: the accessibility of futures that were available before the transition. A political process can close trajectories without making them logically impossible. Laws can be amended, organizations can be rebuilt, expectations can change, and alliances can be reconstructed. The relevant question is practical reachability under bounded time, resources, coordination capacity, and political legitimacy.

Let $\Omega_t(\tau, \rho)$ denote the set of configurations that remain practically reachable from time t within horizon τ and under a resource or disruption bound ρ. The representation is deliberately abstract:

$$\Omega_t(\tau, \rho) := \{x : \ x \text{ is reachable from } x_t \text{ within } \tau \text{ under admissible cost } \rho\}. \tag{19}$$

An intervention, transition, or self-reinforcing sequence can change this set. A possible instance of trajectory closure is represented by

$$\Omega_{t+k}(\tau, \rho) \subsetneq \Omega_t(\tau, \rho), \tag{20}$$

for a substantively meaningful class of alternatives. Equation (20) should not be read as a universal tendency toward shrinking possibility. Transitions can also create new options. The concept of closure refers specifically to lost practical accessibility that is relevant to future collective choice.

Definition 5.5 (Trajectory closure)**.** *Trajectory closure* is a historically situated contraction in the set of practically reachable political configurations produced or reinforced by a transition, where the contraction arises through accumulated commitments, feedback, institutional restructuring, changed capabilities, altered expectations, or other specified mechanisms. Closure can be partial, uneven, and reversible at sufficiently different horizons or resource levels.

The definition contains three safeguards. First, closure is indexed to a horizon and resource bound. A path that is closed over two years can become accessible over two decades. Second, closure can be distributed unevenly. A transition can expand options for one group while sharply narrowing them for another. Third, the relevant possibility space is political rather than purely logical. A technically possible reversal can remain inaccessible when it would require prohibitive coordination, coercion, economic disruption, or violation of legitimate constraints.

This distributional dimension matters for later sections. Aggregate option space can increase while some actors lose meaningful alternatives. Conversely, a measure designed to preserve one group's freedom of action can impose irreversible risks on others. Trajectory closure therefore cannot be evaluated through a system-wide cardinality measure alone. The identity of affected actors, the kinds of options lost, and the legitimacy of the process that closes them remain independent normative questions.

Proposition 5.6 (Closure is relationally distributed)**.** *A transition should be evaluated for the distribution of trajectory closure across actors, institutions, and scales rather than only for aggregate system persistence. Changes that appear reversible at a macro level can coincide with persistent losses of agency or opportunity at a local level.*

5.4 Path dependence and historical sequence

Trajectory closure also clarifies the role of sequence. If the political field were memoryless, the order in which perturbations occurred would have little independent significance once the current state was known. Path-dependent systems violate this simplification. Earlier events can create institutions, expectations, infrastructures, and distributions of capability that alter the effect of later events. Pierson's account of increasing returns in politics highlights how timing and sequence can produce durable consequences and make established courses difficult to reverse (Pierson 2000).

A minimal representation can write the transition kernel as

$$\Pr(x_{t+1} \mid x_t, \eta_t, H_t), \tag{21}$$

where H_t carries trajectory-relevant history that is not reducible to the current observable vector x_t. In practice, H_t can include institutional precedent, sunk investments, remembered commitments, reputational histories, organizational learning, or established expectations. The same visible state can therefore possess different futures when its histories differ.

Near criticality, historical sequence becomes especially significant because a small event can occur at a moment when several trajectories remain accessible. Once one sequence begins to reinforce itself, later events enter a different response field. This is one reason retrospective narratives should remain cautious. A transition that later appears inevitable may have been contingent at the time of branching, while a moment that appeared open may already have contained strong structural asymmetries. Criticality language should therefore increase attention to sequence without turning contingency into retrospective mythology.

Remark 5.7 (Critical junctures and criticality)**.** The language of a political *critical juncture* and the language of a near-critical dynamical region overlap only partially. A critical juncture is a historical-institutionalist concept concerning periods in which choices and sequences can have durable effects. Near criticality additionally proposes a response-sensitive structure in which perturbation, feedback, and transition probability require explicit analysis. The terms should therefore remain distinct unless an empirical study establishes their connection.

5.5 Closure without historical inevitability

The vocabulary of attractors, lock-in, and closure can easily become too strong. Political systems contain interpretation, innovation, coalition change, institutional redesign, exogenous shocks, and strategic agency. A persistent configuration can weaken when actors create new relations or when environmental conditions alter its supporting mechanisms. The existence of self-reinforcement therefore does not establish historical inevitability.

This limitation can be represented by allowing the reachable set itself to change through innovation and relation formation:

$$\Omega_{t+1} = \mathcal{G}(\Omega_t, x_t, G_t, H_t, \eta_t, u_t, \nu_t)\,, \tag{22}$$

where ν_t represents newly generated practices, institutions, coalitions, or interpretive possibilities that were weakly represented or absent from the previous description. Equation (22) prevents option space from being treated as a fixed menu. Political action can generate new possibilities as well as close old ones.

Working Claim 5.8 (Contingent closure)**.** Trajectory closure should be treated as a degree and mechanism rather than as a binary property. A configuration can become substantially harder to leave while remaining transformable through mechanisms that were unavailable, weak, or unrecognized at the moment of entry.

This claim is important for the paper's later engagement with *wuwei* and non-coercive response. Preserving openness does not require preserving every existing option, and allowing a trajectory to develop does not imply passive acceptance of every emerging order. The relevant analytical task is to observe which mechanisms are deepening persistence, which alternatives are becoming costly, and which new possibilities are being generated through the process.

5.6 Transition significance and normative restraint

Attractors, hysteresis, and trajectory closure convert dynamical sensitivity into a problem of political responsibility. Suppose a near-critical configuration permits a small intervention u to raise the probability of entry into a self-reinforcing family B. A purely instrumental analysis can focus on the ratio between intervention cost and response magnitude. The present framework adds a second ratio: the relation between the ease of entry and the difficulty of reopening alternatives after the transition. Schematically,

$$\mathcal{L}_{\text{closure}} := \frac{\text{expected persistence or reversal burden after entry}}{\text{intervention burden at the near-critical state}}. \tag{23}$$

The expression is not proposed as a ready-made empirical index. It identifies a normatively relevant asymmetry: the same near-critical sensitivity that makes steering inexpensive can make the consequences of steering disproportionately persistent.

Proposition 5.9 (Dynamical leverage and trajectory responsibility)**.** *Evidence that a near-critical system can be moved toward a persistent trajectory with a small intervention establishes a potential capability. When the induced trajectory can narrow future options, redistribute reversal burdens, or alter the*

conditions of later collective choice, the same evidence also increases the importance of justification, reversibility analysis, and affected-party consideration.

The proposition deliberately stops short of a decision rule. It does not state that high-leverage interventions should be avoided, because non-intervention can also permit destructive or irreversible transitions. It states that low intervention cost should not be confused with low political consequence. The possibility of return asymmetry and trajectory closure changes the meaning of near-critical information: knowledge of leverage can support restraint as well as action.

Research Obligation 5.10 (Persistence and closure specification)**.** Empirical applications that use attractor, hysteresis, lock-in, or trajectory closure language should identify the mechanisms of persistence, the temporal horizon, the conditions associated with entry and return, the actors whose options change, and the evidence bearing on alternative explanations. Claims of irreversibility should be reserved for cases in which relevant reversal mechanisms have been examined rather than inferred from persistence alone.

The section therefore establishes the temporal asymmetry that motivates the next stage of the paper. A perturbation can reveal a response field; a response can participate in a transition; and a transition can alter the accessibility of later trajectories. Section 6 turns from this dynamical structure to its epistemic problem: how susceptibility, proximity to transition, and noisy signals can be known when the political system is reflexive, historically conditioned, and changing partly in response to attempts at observation.

6 Epistemics of Susceptibility

This section develops the epistemic problem created by the response field of Section 4 and the persistence mechanisms of Section 5. Its role is to ask how susceptibility can be recognized before political actors treat it as a basis for intervention. The discussion has five objectives. First, it distinguishes susceptibility from its observable traces. Second, it develops the idea of *noisy recognition*: near-critical information is commonly partial, model-dependent, delayed, and contaminated by changing perturbation regimes. Third, it separates direct response observation from indirect warning signals such as recovery rates, variance, autocorrelation, spatial propagation, and network response. Fourth, it argues for multi-channel inference rather than single-indicator diagnosis. Fifth, it keeps epistemic recognition separate from intervention warrant. A credible inference that susceptibility has increased can justify attention, monitoring, or restraint without by itself determining a political response.

The scientific literature on critical transitions offers a disciplined starting point for this epistemic problem. In systems approaching some bifurcations, recovery from small perturbations can slow, and statistical indicators such as variance and lagged autocorrelation can change before transition (Scheffer, Bascompte, et al. 2009; Nes and Scheffer 2007). These signals are useful precisely because the underlying loss of resilience is not observed directly. Yet their interpretation is conditional. Reviews of resilience indicators emphasize that critical slowing down is absent in some regime shifts, can occur without bistability, and can be obscured or mimicked by changes in stochastic forcing (Dakos et al. 2015). Quantitative work on early-warning detection likewise shows substantial trade-offs between false alarms and missed detection, even under comparatively favorable assumptions (Boettiger and Hastings 2012). The political lesson is consequently methodological rather than predictive: susceptibility should be inferred through converging evidence with explicit uncertainty, not read directly from one fluctuating indicator.

Figure 5 summarizes the section's epistemic architecture. Susceptibility is treated as a latent property of a situated response field. Observable traces can update beliefs about that property, but none is assumed to be a transparent measurement of it. The distinction matters especially in politics, where observation is sparse, categories are contested, actors learn strategically, and the processes generating observations can change while the observer is still estimating them.

6.1 Latent susceptibility and observable traces

Section 4 represented response through relations between perturbations and subsequent changes. For epistemic purposes, the central object is not the response after it is fully realized but the system's *potential* to respond across possible directions. Let $\chi_t(\eta; h)$ denote a provisional susceptibility profile: the expected magnitude, direction, and propagation of response over horizon h to a perturbation class η, conditional on the current configuration and its history. In a formal model this object may be represented by derivatives, finite-response operators, or network susceptibility matrices. In political analysis, the object is usually only partially identified.

The distinction between susceptibility and observation can be written as

$$\mathcal{O}_t = \mathcal{H}(\chi_t, x_t, G_t, H_t, \eta_t, \epsilon_t)\,, \tag{24}$$

where $\mathcal{O}_t$ is the observed evidence, $\mathcal{H}$ is an observation process, and ϵ_t collects measurement error, omitted variables, and classification uncertainty. Equation (24) does not require a particular statistical estimator. It states that observed fluctuation is mediated by a measurement and interpretation process. The same underlying susceptibility can therefore generate different observable traces under different perturbation regimes or data infrastructures.

Definition 6.1 (Susceptibility inference)**.** *Susceptibility inference* is the provisional estimation of how a situated political configuration may respond to specified perturbation classes, based on observed recov-

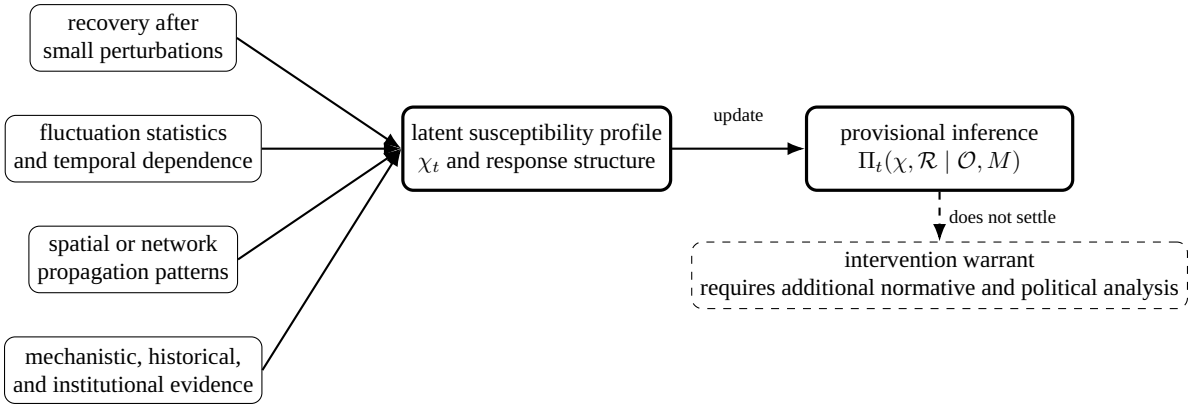

Figure 5: Susceptibility as a latent political property inferred through multiple observational channels. Recovery behavior, fluctuation statistics, network propagation, and mechanistic or historical evidence can jointly inform a provisional susceptibility assessment. The inferential result is separated from intervention warrant because knowledge of sensitivity does not determine the justice, legitimacy, or desirability of exploiting it.

ery, fluctuation, propagation, mechanistic, and historical information, together with explicit uncertainty about the observation and model processes that connect those traces to the latent response structure.

The definition is deliberately plural. A system may be highly susceptible to one perturbation direction and comparatively insensitive to another. A network may transmit economic shocks while buffering informational ones. A political coalition can be stable against gradual pressure but vulnerable to a small coordination event. Treating susceptibility as a single scalar risks erasing precisely the directional and relational structure that matters near a transition.

6.2 Recovery, fluctuation, and indirect recognition

One route to susceptibility is to observe recovery after perturbation. In a simple stable linearization, slower recovery can correspond to an eigenvalue approaching zero in magnitude, reducing the rate at which deviations decay. This logic underlies critical slowing down and motivates early-warning indicators in ecological and other dynamical settings (Scheffer, Bascompte, et al. 2009; Nes and Scheffer 2007). The conceptual point for political analysis is modest: the time structure of recovery can reveal information that the level of a state variable does not.

Suppose a political indicator returns toward a local baseline after a bounded shock. A stylized recovery coefficient can be written as

$$\rho_t(h) := \frac{D(x_{t+h}, x_t^{\text{ref}})}{D(x_t, x_t^{\text{ref}})}, \tag{25}$$

for a distance measure D and a contextually specified reference state. A rise in $\rho_t(h)$ toward one can indicate slower recovery over horizon h. The coefficient is not a universal political early-warning statistic. Its meaning depends on the construction of the reference state, the identity of the perturbation, the stationarity of background conditions, and whether returning toward the reference configuration is substantively meaningful.

When deliberate perturbation is unavailable or ethically inappropriate, natural fluctuations can act as weak probes. Changes in variance, autocorrelation, spatial correlation, or response covariance may contain information about an evolving response field. This is one reason random noise cannot be treated solely as contamination. Section 3 argued that fluctuations can disclose structure. The epistemic refinement here is that the disclosure is indirect: the observer infers latent susceptibility from how fluctuations are received, remembered, and propagated.

Working Claim 6.2 (Fluctuation-probe principle)**.** A fluctuation can be epistemically useful when the system's response to that fluctuation constrains plausible descriptions of susceptibility. The value of the fluctuation lies in the information supplied by the perturbation–response relation, not in randomness by itself.

This principle avoids romanticizing noise. Some fluctuations reveal little. Others are too large to function as local probes. Still others change the system while being observed. The same fluctuation can also contain exogenous forcing and endogenous response components that are difficult to separate. Consequently, recognition depends on repeated comparison, mechanistic context, and uncertainty analysis.

6.3 False alarms, missed alarms, and changing noise regimes

Noisy recognition becomes especially important when warning indicators can be wrong in both directions. Boettiger and Hastings demonstrate that commonly used early-warning statistics can exhibit substantial false-positive and false-negative rates, motivating explicit characterization of reliability and sensitivity (Boettiger and Hastings 2012). Dakos and colleagues similarly show that changes in stochastic forcing can generate apparent warning trends, while large process or observation noise can hide a genuine loss of resilience (Dakos et al. 2015). These findings block a simple inference from "more fluctuation" to "closer to transition."

For political analysis, the relevant error structure can be represented by a fourfold distinction:

	transition-relevant susceptibility	no such susceptibility
warning recognized	informative alarm	false alarm
warning not recognized	missed alarm	quiet observation

(26)

The labels are epistemic rather than moral. A false alarm can still reveal a real disturbance regime, and a quiet observation does not prove robust stability. The table instead clarifies that warning systems face asymmetric risks whose political costs can differ sharply.

Definition 6.3 (Noisy recognition)**.** *Noisy recognition* is an epistemic condition in which observations alter the rational assessment of susceptibility without uniquely determining whether a critical transition is approaching, which mechanism is responsible, or what response should follow.

Noisy recognition is not a temporary defect expected to disappear with enough data. Better data can reduce uncertainty, but political systems contain structural sources of ambiguity: changing actor strategies, heterogeneous time scales, incomplete observation, endogenous institutions, and shifting classification practices. Near-critical inquiry should therefore represent uncertainty as part of the object rather than as an embarrassment around an otherwise exact estimate.

Proposition 6.4 (Warning asymmetry)**.** *The presence of a credible warning signal can increase concern about elevated susceptibility, while the absence of such a signal does not generally establish low susceptibility. This asymmetry is especially relevant when the detection process is known to contain missed alarms, transition mechanisms that do not produce the selected indicator, or observation noise capable of masking the signal.*

The proposition is deliberately weaker than a precautionary decision rule. It concerns inference only. Whether the asymmetry should lead to monitoring, buffering, public communication, restraint, or intervention remains a separate question.

6.4 Model dependence and historical baselines

Warning statistics acquire meaning relative to a model and a baseline. A rise in variance can represent declining restorative strength, stronger exogenous shocks, altered measurement practices, changing population composition, or several processes at once. An increase in temporal dependence can reflect critical

slowing down, institutional memory, synchronized information flows, or a change in the time scale of external forcing. The same numerical trend therefore supports different interpretations under different causal models.

Let $M_1, \ldots, M_K$ denote competing descriptions of the response process. A susceptibility assessment is better represented by

$$\Pi_t(\chi, \mathcal{R}) = \sum_{k=1}^{K} \Pi_t(\chi, \mathcal{R} \mid M_k)\, w_t(M_k), \tag{27}$$

where $\mathcal{R}$ denotes a response structure and $w_t(M_k)$ represents the current evidential weight assigned to model M_k. Equation (27) is schematic rather than a demand for Bayesian implementation. Its purpose is to keep model uncertainty visible.

Historical baselines create a related problem. Early-warning research notes that relative changes from a baseline can sometimes be more informative than a single measurement of an indicator (Dakos et al. 2015). In politics, however, the baseline itself can be historically unstable. Institutions, technologies, demographic composition, communication systems, and actor expectations can change the response field. The baseline should therefore be treated as a dated comparison condition rather than an ahistorical normal state.

Remark 6.5 (Baseline historicity)**.** A political baseline is part of the historical configuration being studied. Evidence that a present fluctuation differs from an earlier baseline can be informative while still leaving open whether the difference reflects changing susceptibility, changing forcing, changing measurement, or a new political regime.

This point becomes especially important for international politics, where a long time series can span transformations in institutions, alliance structures, technologies, and communication speed. Statistical continuity of an indicator does not guarantee continuity of the mechanism generating it.

6.5 Triangulation across response channels

The limitations of individual indicators motivate a triangulated epistemic strategy. The aim is not to combine weak signals mechanically into a strong score. It is to ask whether different channels constrain the same susceptibility hypothesis. Useful channels can include:

- recovery behavior after comparable small disturbances;
- trends in variance, autocorrelation, or other fluctuation statistics;
- changes in spatial, organizational, or network propagation;
- mechanism-specific evidence concerning feedback and restorative forces;
- historical evidence concerning previous entry, recovery, and persistence;
- cross-scale observations showing whether local deviations remain local or recruit wider structures.

These channels need not agree. Disagreement can itself be informative. Rising variance with stable recovery may suggest changing perturbation intensity rather than declining restorative strength. Local fragility with weak network propagation may imply compartmentalized susceptibility. Strong propagation without slowing recovery may indicate a coupling problem rather than proximity to a local bifurcation. Near-critical epistemics therefore benefits from preserving heterogeneity among indicators long enough to diagnose why they differ.

Working Claim 6.6 (Triangulation without forced convergence)**.** A susceptibility assessment becomes more informative when multiple independent or differently biased channels constrain the same response hypothesis. Lack of convergence should remain visible when channels imply different mechanisms, rather than being hidden by aggregation into a single warning score.

The claim also limits the use of dashboards. A composite index can be convenient for communication, but it can destroy the structural information needed to understand whether susceptibility is localized, directional, delayed, or network-mediated. Near criticality, preserving the architecture of evidence can matter more than reducing it to one number.

6.6 Passive observation, active probing, and epistemic intervention

Some forms of susceptibility can be estimated more directly by perturbing a system and observing its response. In laboratory and engineered settings, such probing can be routine. Political systems impose stronger constraints. A probe can affect rights, expectations, material welfare, and strategic behavior. It can also trigger the transition whose proximity the observer hoped to measure. The epistemic value of perturbation must therefore be separated from permission to perform it.

This separation is particularly important near criticality because a small test can have a large response. An experiment designed to estimate a local derivative may cross a practical basin boundary. The existence of potentially high information value does not remove the ethical burden created by that possibility. Accordingly, this paper treats naturally occurring fluctuations, historical variation, simulations, and low-consequence institutional experiments as different evidence classes rather than interchangeable sources.

Research Obligation 6.7 (Epistemic intervention specification)**.** When susceptibility estimates depend on deliberate perturbation, the analysis should specify the perturbation's expected scale, reversibility, affected parties, possible propagation channels, stopping conditions, and the reasons the information could not be obtained through lower-consequence observation. The resulting evidence should be reported separately from the normative justification for conducting the probe.

This obligation anticipates the paper's later sections on justice and irreversibility. It also marks a limit to the analogy with physical response theory. Political observation can be intervention, and intervention can alter the response field that observation seeks to estimate.

6.7 Epistemic outputs and action restraint

The output of near-critical epistemics should therefore be richer than a binary classification such as "critical" or "safe." A provisional assessment can instead report a structured object

$$\mathcal{E}_t = \left(\widehat{\chi}_t, \widehat{\mathcal{R}}_t, \mathcal{U}_t, \mathcal{A}_t, \mathcal{M}_t\right), \tag{28}$$

where $\widehat{\chi}_t$ is an estimated directional susceptibility profile, $\widehat{\mathcal{R}}_t$ is a provisional response architecture, $\mathcal{U}_t$ records uncertainty, $\mathcal{A}_t$ records known ambiguities or alternative mechanisms, and $\mathcal{M}_t$ records the measurement and model conditions under which the assessment was produced.

Proposition 6.8 (Epistemic recognition and intervention warrant)**.** *Recognition of elevated susceptibility can strengthen the case for attention, monitoring, option preservation, and analysis of irreversible risk. It does not by itself determine whether a perturbation should be suppressed, amplified, redirected, or deliberately introduced. The transition from susceptibility knowledge to political action requires additional judgments concerning legitimacy, distribution, justice, ethics, and the value of alternative trajectories.*

The proposition is central to the paper's departure from control-theoretic reasoning. Near-critical information can reveal leverage, but leverage is not a normative instruction. Indeed, the more conse-

quential a small perturbation may be, the more important it becomes to distinguish epistemic capability from intervention entitlement.

Research Obligation 6.9 (Uncertainty-preserving reporting)**.** Empirical claims of near-critical susceptibility should report the perturbation classes considered, temporal horizon, observation process, relevant baselines, competing mechanisms, indicator disagreements, and plausible false-alarm and missed-alarm pathways. A single point estimate or warning label should not substitute for this structure when the omitted uncertainty is relevant to political response.

The section therefore replaces the ideal of exact recognition with a more limited epistemic objective: identify how evidence changes the plausible response field while preserving uncertainty about mechanism, proximity, and future transition. This prepares the argument of Section 7. Tangent-space reasoning offers one disciplined way to use local information under global ignorance, but near critical regions the relation between local validity and intervention safety becomes unstable. The next section develops that boundary.

7 Tangent-Space Reasoning near Critical Regions

This section develops a boundary condition for tangent-space reasoning. Its role is to examine what happens when locally credible information is available in a configuration whose response field is becoming highly susceptible. The section has five objectives. First, it defines tangent-space reasoning as a disciplined use of local directional information under global ignorance. Second, it separates three neighborhoods that are easily conflated in policy analysis: an epistemic neighborhood in which a local approximation remains informative, a dynamical neighborhood in which intervention remains sufficiently reversible, and a normative neighborhood in which intentional intervention has adequate political warrant. Third, it shows why these neighborhoods can diverge near critical regions. Fourth, it argues that susceptibility information can acquire an inverted practical meaning: improved knowledge of leverage can strengthen reasons for restraint alongside reasons for active use. Fifth, it reformulates tangent-space reasoning as a method for locating limits, dangerous directions, and conditions for further inquiry, with action authorization treated as a separate political judgment.

The physical literature provides a useful contrast. Local response methods can identify how collective dynamics change under small parameter variations, and work on network susceptibilities formalizes directional responses of nodes and edges near steady states (Manik et al. 2017). Classical chaos-control results go further: under specified dynamical conditions, very small parameter perturbations can stabilize unstable periodic motions or direct trajectories toward accessible targets (Ott et al. 1990; Shinbrot et al. 1990). These results demonstrate that local sensitivity can create substantial leverage. Transfer to political analysis would require additional grounds for observability, controllability, stable parameterization, and legitimate target selection. The political problem developed here begins precisely where the engineering attraction of leverage becomes normatively and epistemically ambiguous.

Figure 6 summarizes the central distinction. The figure records a conceptual separation among three questions; empirical applications would require context-specific estimates of their boundaries: where local inference is credible, where intervention remains dynamically revisable, and where intentional action is politically justified. Near criticality, treating these questions as one can turn an epistemic advantage into a governance error.

7.1 Local directional reasoning under global ignorance

Let a political configuration be represented abstractly by a state x and a set of intervention variables u. Around a current configuration (x_t, u_t), a smooth model can be written locally as

$$\delta\dot{x} = J_x(x_t, u_t)\,\delta x + J_u(x_t, u_t)\,\delta u + \mathcal{O}\big(\|\delta x\|^2 + \|\delta u\|^2\big)\,, \tag{29}$$

where J_x and J_u summarize local directional response. Equation (29) should be read as a formal sketch. Political analysis rarely observes a full state vector or a stable Jacobian, and the relevant variables can change during the period of interest. The value of the expression is conceptual: it separates knowledge of nearby directions from knowledge of the full global landscape.

Definition 7.1 (Tangent-space reasoning)**.** *Tangent-space reasoning* is the use of locally supported directional information to characterize nearby response possibilities while restricting claims to the scope supported by available knowledge of the local configuration, perturbation class, and time horizon. In political application, it includes explicit limits on horizon, perturbation class, model stability, and re-estimation conditions.

The attraction of this approach is epistemic modesty. A policymaker may possess credible evidence that one variable responds positively to a small adjustment while possessing only partial knowledge of the long-run future. Under comparatively regular conditions, bounded intervention followed by renewed observation can be more defensible than global optimization. Complexity-governance research similarly

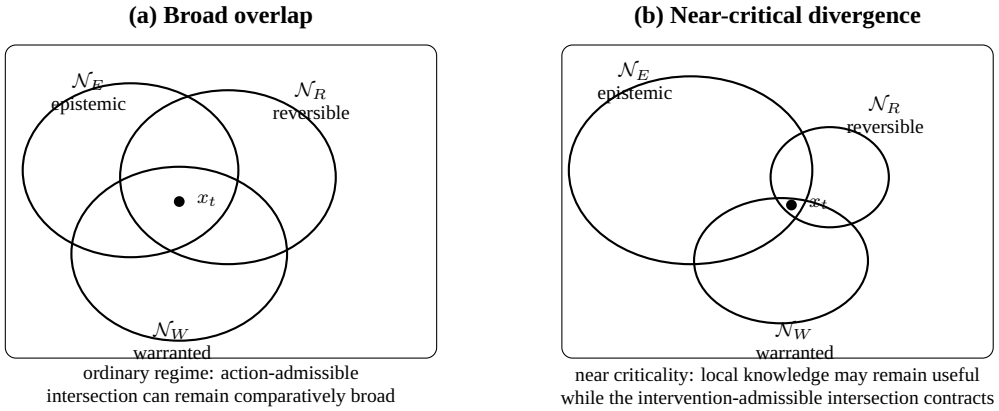

Figure 6: Conceptual separation of three local neighborhoods. $\mathcal{N}_E$ denotes a region in which a local model remains epistemically informative; $\mathcal{N}_R$ denotes a region in which intervention remains sufficiently reversible relative to specified horizons and loss criteria; and $\mathcal{N}_W$ denotes a region of normatively warranted intervention. The sets can overlap in several configurations. Near criticality, their intersection can become small even while local directional information remains available.

emphasizes adaptation, feedback, and limited predictability while treating governance actors as participants within the system (Duit and Galaz 2008; May 2022).

Near critical regions, however, the meaning of "local" changes. A small movement in state space can carry the system across a basin boundary, alter a feedback architecture, or place subsequent dynamics under a different regime. The local approximation may remain informative about the immediate direction of motion while becoming weak evidence about the consequences of acting on that direction.

7.2 Epistemic, reversible, and warranted neighborhoods

The distinction can be formalized through three provisional neighborhoods. Let $\mathcal{N}_E(x_t)$ denote the set of nearby configurations and perturbations for which a local representation remains sufficiently informative for a specified inferential purpose. The tolerance can concern prediction error, qualitative direction, ranking of responses, or another task-specific criterion. Let $\mathcal{N}_R(x_t)$ denote the set for which intervention is expected to remain sufficiently reversible over horizon h relative to specified loss and option criteria. Finally, let $\mathcal{N}_W(x_t)$ denote the set of interventions for which there is adequate normative and political warrant given affected parties, rights, distribution, authority, and procedural conditions.

The locally admissible action region can then be represented as

$$\mathcal{N}_A(x_t) = \mathcal{N}_E(x_t) \cap \mathcal{N}_R(x_t) \cap \mathcal{N}_W(x_t). \tag{30}$$

Equation (30) expresses a conjunction of conditions for action admissibility. Epistemic validity, dynamical reversibility, and normative warrant are separate necessary considerations for intentional trajectory-shaping. In some contexts one of the sets may be difficult to estimate. That difficulty should remain visible through explicit uncertainty about the relation among the sets.

Proposition 7.2 (Neighborhood divergence near criticality)**.** *As susceptibility rises and persistent transition becomes more plausible, an epistemically credible local neighborhood can remain broader than the region in which intervention is dynamically reversible or normatively warranted. Local knowledge can therefore survive a contraction of the action-admissible neighborhood.*

The proposition is conditional, with scope varying across transition types and political contexts. Some critical transitions are reversible, some interventions remain low consequence, and some settings supply strong prior justification for action. The analytical point is that the three neighborhoods have

different determinants. Improved estimation of J_u can expand $\mathcal{N}_E$ while information about hysteresis or trajectory closure contracts $\mathcal{N}_R$. A clear emergency mandate can expand part of $\mathcal{N}_W$ even when epistemic uncertainty remains high. The framework prevents any one dimension from silently substituting for the others.

7.3 Local validity and trajectory discontinuity

Near a critical region, the immediate response to a perturbation can be smooth while the long-run trajectory changes discontinuously. This creates a horizon problem. Let $\Phi_h(x_t, u)$ denote the state reached after horizon h. A small intervention may satisfy

$$\|u_1 - u_0\| \ll 1 \qquad \text{and} \qquad \|\Phi_{h_0}(x_t, u_1) - \Phi_{h_0}(x_t, u_0)\| \ll 1 \tag{31}$$

for a short horizon h_0, while producing a much larger separation over a longer horizon h_1 because the two trajectories enter different response regimes. Noise-induced transition theory supplies one physical illustration of how stochastic perturbations can produce state changes that are poorly represented by deterministic local intuition (Forgoston and Moore 2018). Critical-transition research likewise emphasizes that loss of recovery capacity can precede abrupt regime change (Scheffer, Bascompte, et al. 2009; Nes and Scheffer 2007).

The political implication is a distinction between *local model validity* and *trajectory continuity*. A model can correctly estimate an immediate marginal effect while remaining inadequate for the politically relevant time horizon. This is particularly consequential when the transition activates increasing returns, institutional reinforcement, or path dependence that raises the later cost of reversal (Pierson 2000). Section 5 called this contraction of practically reachable alternatives trajectory closure.

Working Claim 7.3 (Local-effect sufficiency)**.** Evidence supporting a reliable short-horizon marginal effect supplies only one component of near-critical decision analysis. When persistent transition is plausible, the analysis should also examine whether the intervention changes regime membership, feedback architecture, or the practical cost of returning to previously reachable configurations.

The working claim preserves the value of marginal evidence while limiting its scope. A local derivative can describe a nearby response; continuity of the future landscape requires additional analysis.

7.4 Critical susceptibility and the inversion of leverage

High susceptibility often appears attractive from a control perspective because small inputs can produce large effects. The classical chaos-control literature illustrates the engineering form of this attraction: sensitivity can be used to stabilize or redirect trajectories with small perturbations when the relevant structure is sufficiently characterized (Ott et al. 1990; Shinbrot et al. 1990). In politics, the same fact has a second interpretation. If a small intervention can reorganize a large collective trajectory, then error, bias, illegitimate authority, and distributive asymmetry can also be amplified.

This yields an inversion of the ordinary relation between knowledge and action. Away from a sensitive boundary, better information about a local effect often supports more confident bounded intervention. Near a persistent transition, better information that reveals unusually high leverage can instead increase the responsibility attached to acting. The relevant practical inference can be monitoring, protection of reversibility, consultation, or restraint.

Proposition 7.4 (Leverage–restraint inversion)**.** *Under conditions of high susceptibility, persistent consequence, and material normative uncertainty, improved evidence that identifies a high-leverage direction can strengthen reasons for restraint as well as reasons for action. The direction of practical implication depends on reversibility, legitimacy, distribution, and the consequences of error, with leverage forming one component of that judgment.*

The proposition is central to the paper's later analysis of evidence-based policy. Evidence that identifies a sensitive parameter is valuable. Near criticality, that value can take the form of monitoring, protection against accidental forcing, preservation of option space, or identification of a region in which further intervention requires a higher justificatory threshold.

7.5 Performative local knowledge in reflexive systems

Political tangent-space reasoning also differs from local reasoning in many physical systems because actors can observe, interpret, and strategically respond to the analysis itself. A public warning that a coalition is fragile, a market is near a threshold, or an institutional settlement is vulnerable can alter expectations and behavior. The estimated response field can therefore enter the dynamics it describes.

A minimal representation is

$$x_{t+1} = F\Big(x_t, u_t, \eta_t, \mathcal{I}_t(\widehat{J}_t)\Big)\,, \tag{32}$$

where $\widehat{J}_t$ denotes a local susceptibility estimate and $\mathcal{I}_t(\widehat{J}_t)$ represents the institutional and strategic effects of communicating, withholding, or acting upon that estimate. The expression makes a limited point: local information can change the response field through its circulation.

This reflexivity adds a fourth form of uncertainty to the three neighborhoods of Figure 6. The analyst must consider whether the estimate is correct, whether intervention is reversible, whether it is warranted, and how production and disclosure of the estimate changes the object of inference. Complexity-governance scholarship's emphasis on actors embedded within adaptive systems is relevant here (Duit and Galaz 2008; May 2022).

Research Obligation 7.5 (Reflexive-information specification)**.** When near-critical analysis is expected to alter strategic behavior, the assessment should specify plausible channels through which publication, classification, warning, secrecy, or policy uptake can change the response field. The performative pathway should be treated as part of both intervention analysis and communication analysis.

This obligation is especially important when a warning can coordinate actors in the same direction, generate anticipatory withdrawal, or create incentives to exploit a revealed vulnerability. In such cases, evidence is simultaneously an observation and a potential perturbation.

7.6 Boundary-aware tangent-space reasoning

The limitations developed above preserve local reasoning while assigning it a different function. Global knowledge remains unavailable in many political settings, and waiting for a full map of the landscape can be as irresponsible as acting on a narrow derivative. Near critical regions, tangent information can serve primarily to locate sensitivities, boundaries, and conditions for further inquiry.

Instead of using local information primarily to select the most effective steering direction, a criticality-aware procedure can use it to identify: (i) directions of elevated susceptibility; (ii) candidate basin boundaries or feedback changes; (iii) variables whose perturbation should be treated as high-consequence; (iv) evidence needed to estimate reversibility; (v) actors whose options or burdens may change disproportionately; and (vi) conditions under which further action should pause for renewed observation or political authorization.

Definition 7.6 (Boundary-aware tangent-space reasoning)**.** *Boundary-aware tangent-space reasoning* is the use of local directional information together with explicit estimates of transition proximity, reversibility, model horizon, reflexive effects, and normative warrant. Its purpose is to characterize locally available possibilities and limits while preserving the distinction between knowing a sensitive direction and possessing permission to exploit it.

A practical output can therefore be represented as

$$\mathcal{T}_t = \left(\widehat{J}_t, \mathcal{N}_E, \mathcal{N}_R, \mathcal{N}_W, \mathcal{B}_t, \mathcal{P}_t\right), \tag{33}$$

where $\mathcal{B}_t$ records suspected transition boundaries and $\mathcal{P}_t$ records performative pathways associated with the analysis and its use. The object is deliberately richer than a recommended control vector. It can support action, restraint, monitoring, or institutional consultation depending on the surrounding conditions.

Research Obligation 7.7 (Criticality boundary of local action)**.** An analysis that recommends intentional intervention from local evidence should state the horizon over which local validity is claimed, the evidence concerning trajectory continuity and reversibility, the estimated relation between the intervention and possible regime change, the parties exposed to persistent loss, and the political basis for treating the intervention as warranted. Where these elements are materially unresolved, local efficacy should be reported separately from action recommendation.

The obligation supplies the central boundary condition of this section. Near criticality, epistemic capability and intervention entitlement can separate. The separation is a constitutive feature of political action under nonlinear sensitivity, uncertain irreversibility, and contested value. Additional models can refine each component while leaving their conceptual distinction intact.

7.7 Transition to governance within the perturbation field

Tangent-space reasoning began from a local observer attempting to act under global ignorance. Near criticality, that picture must be widened. The observer, policy instrument, public warning, institutional delay, and decision to refrain from acting all enter a field already populated by endogenous fluctuations, strategic actors, and cross-scale feedback. The next section therefore moves from the validity of local reasoning to the location of governance itself. Rather than positioning governance outside the perturbation field, it treats intentional political response as one perturbation-generating and perturbation-absorbing process among others. This shift is necessary before the paper can develop a repertoire that includes buffering, redirection, selective amplification, restraint, and non-coercive response while treating the governor as a situated participant within the perturbation field.

8 Governance within the Perturbation Field

This section relocates governance from an external steering position to a situated position within the perturbation field. Its role is to examine how an intentional political response enters a system whose actors, expectations, feedback structures, and endogenous fluctuations are already evolving. The section has five objectives. First, it defines governance as a participant in the response field rather than as an analytically separate control layer. Second, it distinguishes the direct effect of an intervention from the secondary perturbations generated through interpretation, anticipation, institutional adjustment, and strategic reaction. Third, it develops the idea of a governance footprint that records how intervention changes both the current configuration and the conditions of subsequent response. Fourth, it clarifies why restraint, delay, and selective non-intervention also have dynamical effects and should therefore be analyzed within the same field. Fifth, it establishes a reflexive criterion for near-critical response: an intervention should be evaluated partly through the perturbations and feedbacks that its own existence is likely to generate.

Complexity-oriented governance research has long emphasized nonlinear change, threshold effects, cascades, limited predictability, and the embeddedness of governance arrangements within complex adaptive systems (Duit and Galaz 2008; May 2022). Policy-feedback research adds a complementary temporal insight: policy can become a cause of subsequent politics by changing resources, incentives, interpretations, and organizational conditions (Pierson 1993). These literatures do not establish that political systems instantiate one universal critical mechanism. They do, however, support a conceptual shift that is central here. Political response is part of the environment to which later political actors respond.

Figure 7 summarizes the section's main relocation. The intervention u_t is analytically distinguishable from an endogenous fluctuation η_t because it is intentionally selected by an actor or institution. Once enacted, however, the intervention becomes an event inside the system. Other actors can observe it, infer intentions from it, reallocate resources, coordinate, resist, imitate, delay, or revise their own expectations. The intervention therefore becomes part of the causal and interpretive environment from which the next political configuration emerges.

8.1 Embedded location of political response

A simple control representation often places a decision maker outside the system:

$$x_{t+1} = F(x_t, \eta_t) + B_t u_t. \tag{34}$$

Equation (34) can be useful for limited analytical purposes, but its architecture can encourage a misleading interpretation in political settings. It separates the system dynamics F from the intervention u_t and can thereby suggest that the intervention enters without changing the rules, expectations, or relations that generate later dynamics.

A more appropriate formal sketch for reflexive political systems is

$$x_{t+1} = F_t(x_t, \eta_t, u_t, \mathcal{R}_t(u_t), \mathcal{E}_t(u_t)), \tag{35}$$

where $\mathcal{R}_t(u_t)$ represents strategic and organizational responses to the intervention and $\mathcal{E}_t(u_t)$ represents expectation and interpretive effects. The time index on F_t is deliberate. Political intervention can change the configuration within which later relations are formed, so the relevant response map can itself evolve.

Definition 8.1 (Embedded governance)**.** *Embedded governance* denotes intentional political response undertaken by actors who remain participants in the system whose trajectory they seek to influence. Their interventions enter the same field of interpretation, feedback, resource redistribution, and strategic adaptation that they attempt to govern.

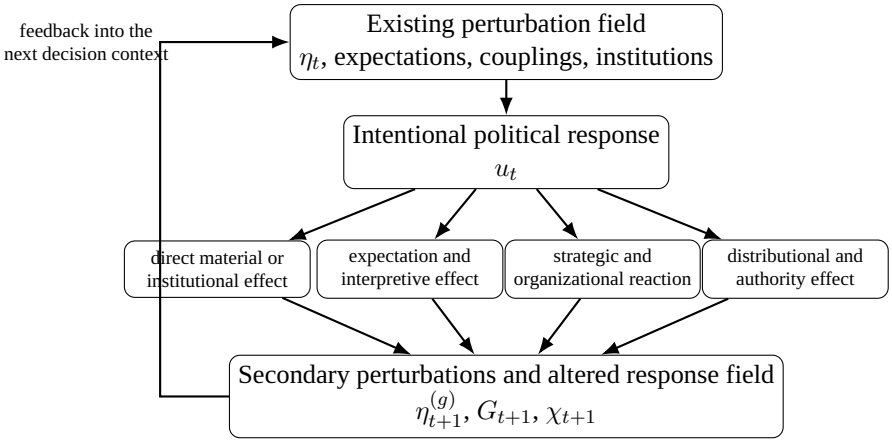

Figure 7: Governance located within the perturbation field. An intentional response has a direct effect and can also alter expectations, strategies, organizational conditions, distribution, and perceived authority. These channels generate secondary perturbations and can modify the response field faced by subsequent actors. The diagram is conceptual: the listed channels need not appear with equal strength in every political setting.

The definition does not imply that every political action has system-wide consequences. Many interventions remain local, weakly coupled, or quickly absorbed. Its purpose is to prevent the analytical position of the policymaker from being confused with an ontological position outside the system. The embedded actor can possess authority and asymmetric resources while remaining subject to feedback generated by the intervention itself.

This distinction matters more as susceptibility rises. In a comparatively regular regime, the secondary effects of a policy can remain small relative to its direct effect. Near a sensitive boundary, a secondary expectation shift or strategic reaction can be amplified through the same couplings that made the original system near critical. An intervention selected because it appears small in instrument space can therefore have a large footprint in response space.

8.2 Direct effects and response-generated perturbations

The perturbation generated by governance can be separated conceptually into a direct component and a response-generated component. Let

$$\Delta_t^{(u)} = \Delta_t^{\text{dir}}(u_t) + \Delta_t^{\text{resp}}(u_t; x_t, G_t, H_t), \tag{36}$$

where G_t denotes relevant coupling structures and H_t denotes historically accumulated relations, commitments, and expectations. The first term records the intended or immediate institutional effect. The second records changes produced because actors and organizations respond to the intervention as an event with political meaning.

Definition 8.2 (Response-generated perturbation)**.** A *response-generated perturbation* is a change in the political field that arises through actors' reactions to an intervention, including revised expectations, coordination, resistance, imitation, strategic repositioning, organizational adaptation, and reinterpretation of institutional commitments.

Policy-feedback scholarship provides one established route through which such effects occur. Policies can alter resources and incentives while also providing information and cues that affect political interpretation (Pierson 1993). The present framework broadens the object of interest from long-run policy feedback to near-critical response. When susceptibility is high, the temporal distance between

intervention and feedback can shrink, and feedback can participate in the transition rather than merely shape politics after a stable policy is established.

This gives a second meaning to the claim that governance can become noise. The term does not erase intentionality. It emphasizes that, from the position of other interacting actors, a governmental decision can arrive as an externally introduced disturbance whose downstream effects are not fully controlled by its initiator. The same intervention can be interpreted differently across subsystems, creating heterogeneous local perturbations that later interact.

Proposition 8.3 (Intervention–perturbation duality)**.** *An intentional intervention can simultaneously function as a selected policy instrument for its initiator and as a perturbation for other actors and subsystems. Near criticality, the political significance of the intervention depends partly on these secondary response pathways rather than solely on its direct intended effect.*

The proposition adds a systems-level caution to instrument evaluation. A policy can perform well relative to its direct target while generating feedback that changes the surrounding response field. Evaluation near criticality should therefore consider both target efficacy and field effects.

8.3 Expectation channels and anticipatory response

Political actors frequently respond to anticipated consequences rather than to completed consequences. An intervention can therefore influence trajectories before its material effects have fully appeared. Announcements, classifications, public warnings, emergency declarations, institutional commitments, and visible preparations can change beliefs about what other actors are likely to do. Those beliefs can alter present behavior.

A minimal representation is

$$a_{i,t+1} = \mathcal{A}_i(x_t, u_t, \mathbb{E}_i[x_{t+h} \mid u_t, \mathcal{I}_{i,t}]) \,, \tag{37}$$

where $a_{i,t+1}$ denotes actor i's subsequent action and $\mathcal{I}_{i,t}$ denotes the information and interpretive resources available to that actor. Equation (37) does not assume rational-expectations equilibrium. It records only that expected futures can mediate present response.

Near criticality, expectation channels can become especially consequential for two reasons. First, many actors may observe the same intervention and update in correlated directions, producing synchronization. Second, uncertainty about the system can make the intervention itself a high-weight signal. A large or unusual policy response can be interpreted as evidence that decision makers possess private information about danger, which can generate further reactions independent of the policy's direct effect.

This possibility complicates public communication. Greater transparency can support accountability and coordination while also changing the response field. Secrecy can reduce immediate coordination around a signal while weakening legitimacy, distributed learning, or later trust. The framework therefore does not produce a universal communication rule. It identifies communication design as part of the perturbation analysis.

Research Obligation 8.4 (Expectation-channel analysis)**.** When an intervention is likely to be interpreted as information about the system, near-critical analysis should examine plausible expectation pathways in addition to direct material effects. The assessment should identify who can observe the intervention, what information they may infer from it, whether responses are likely to be correlated, and whether those responses can alter transition risk.

The obligation follows directly from the embedded location of governance. A policy maker cannot evaluate the intervention solely as a physical input while other actors treat it as a signal, commitment, threat, concession, or forecast.

8.4 Governance cascades and endogenous amplification

A response-generated perturbation can itself become the input to a subsequent round of action. This creates the possibility of a governance cascade. Let $u_t^{(0)}$ denote an initial intervention. A stylized sequence can be written as

$$\eta_{t+1}^{(1)} = H_1(u_t^{(0)}, x_t), \tag{38}$$

$$u_{t+1}^{(1)} = K_1(\eta_{t+1}^{(1)}, x_{t+1}), \tag{39}$$

$$\eta_{t+2}^{(2)} = H_2(u_{t+1}^{(1)}, x_{t+1}), \tag{40}$$

where H_k maps interventions into response-generated perturbations and K_k represents subsequent intentional reactions. The indices do not denote a fixed hierarchy. They mark successive rounds in which response and intervention can alternate.

Under ordinary conditions, the sequence may damp quickly. Near criticality, positive feedback can amplify successive rounds. One institution responds to a fluctuation; other actors interpret the response as evidence of elevated risk; they change behavior; the first institution interprets those changes as confirmation; further intervention follows. The resulting trajectory can be partly generated by the feedback between governance and response.

Working Claim 8.5 (Governance-cascade possibility)**.** When intervention effects are strongly coupled to expectations and subsequent policy responses, near-critical analysis should treat the governance sequence itself as a possible amplification pathway. The relevant object is the interaction between original disturbance, intentional response, and the secondary perturbations generated by that response.

The claim should not be read as an argument for passivity. A timely intervention can interrupt a destructive cascade, decouple subsystems, or create a buffer. The analytical consequence is that intervention should be assessed recursively: what response does it generate, and what response is likely to follow from that response?

8.5 Restraint, delay, and the dynamics of non-intervention

The embedded perspective also changes the analysis of restraint. A decision to delay, refrain, maintain an existing rule, or decline to amplify a perturbation can affect expectations and resource allocation. Non-intervention therefore cannot be treated as a zero vector in every political setting.

Let $u_t = 0$ denote absence of a new active instrument. The system can still evolve according to

$$x_{t+1} = F_t(x_t, \eta_t, \mathcal{R}_t(0), \mathcal{E}_t(0))\,. \tag{41}$$

Here $\mathcal{R}_t(0)$ and $\mathcal{E}_t(0)$ record responses to the absence of new intervention where that absence is politically observable and interpretable. Silence can be read as confidence, incapacity, tolerance, uncertainty, or strategic patience depending on context. The interpretation can itself alter behavior.

This point prevents a false opposition between intervention and wuwei that will be developed later in the paper. Responsive restraint is an active relation to the evolving configuration. Its quality depends on what is being allowed to continue, what boundaries remain protected, and what learning becomes possible during the interval of non-forcing.

Remark 8.6 (Null intervention)**.** A null intervention has low direct instrumental magnitude by construction, but its political footprint can still be nonzero when actors interpret the absence of action and alter behavior accordingly. Near-critical analysis should therefore compare active intervention with plausible trajectories under observable restraint rather than treating restraint as absence of dynamics.

The remark also clarifies the ethical stakes. Choosing restraint can preserve option space in one setting and expose vulnerable actors to accumulating harm in another. A response-centered framework therefore evaluates the relational and distributional consequences of restraint alongside those of active steering.

8.6 Governance footprints and recursive evaluation

The preceding arguments can be consolidated through the concept of a governance footprint. The footprint is broader than implementation output. It records the set of direct and secondary changes through which an intervention modifies the field of subsequent response.

Definition 8.7 (Governance footprint)**.** For an intervention u_t, the *governance footprint* is the provisional set

$$\Phi_t(u_t) = \{\Delta x_t, \Delta G_t, \Delta \chi_t, \Delta \mathbb{E}_t, \Delta \mathcal{D}_t, \Delta \Omega_t\}, \tag{42}$$

where the components denote changes in configuration, coupling, susceptibility, expectations, distributional conditions, and reachable option space. The set is an analytical checklist rather than a claim that each component is measurable on a common scale.

The footprint helps separate a narrow policy-success question from a near-critical response question. A policy may achieve its immediate target while increasing susceptibility elsewhere. It may reduce one source of noise while synchronizing previously independent subsystems. It may stabilize a present configuration while closing future alternatives. It may also create valuable institutional learning, increase buffering capacity, or open a path for decentralized adaptation. Each possibility belongs to the response field.

Recursive evaluation therefore asks at least three temporal questions. First, what does the intervention change directly? Second, what perturbations and expectations does that change generate? Third, how do those secondary effects alter the conditions under which the next intervention will be interpreted and implemented? Policy-feedback research supports this temporal orientation by treating policy effects as possible causes of later political development (Pierson 1993), while path-dependence research emphasizes how self-reinforcement can make later reversal increasingly difficult (Pierson 2000).

Proposition 8.8 (Recursive intervention criterion)**.** *Near criticality, evaluation of an intentional response should extend beyond its direct target effect to the response-generated perturbations and field changes that materially alter subsequent susceptibility, coupling, distribution, legitimacy, or option space. The stronger those recursive effects are expected to be, the weaker the case for treating intervention as an isolated policy instrument.*

This criterion also links back to Section 7. A locally credible marginal effect can remain valid while the governance footprint extends well beyond the neighborhood in which the direct effect was estimated. The relevant question is therefore not simply whether the intervention works near x_t, but whether the intervention changes what "near x_t" will mean for the next round of political response.

8.7 Transition to response repertoires

Locating governance within the perturbation field changes the practical problem. The task is no longer exhausted by choosing between intervention and non-intervention. Political systems possess multiple ways of receiving and transforming perturbations, and intentional actors can participate in those responses with different degrees of force, delay, reversibility, and coordination. The next section therefore develops response repertoires and adaptive capacity. It will distinguish absorption, buffering, decoupling, redirection, selective amplification, delay, experimentation, and restraint as possible modes whose appropriateness depends on the perturbation–response relation and the normative conditions of the case.

9 Response Repertoires and Adaptive Capacity

The preceding sections shifted the analytical focus from isolated perturbations toward relations between perturbation and response, and then located intentional governance inside the same evolving field. This section develops the next consequence of that shift. Near criticality, adaptive capacity should not be identified with the strength of a single control instrument. It depends partly on whether a political system can receive heterogeneous disturbances through a sufficiently varied repertoire of responses while preserving learning, reversibility, and normative scrutiny. The structural role of the section is to specify that repertoire, distinguish response diversity from indiscriminate flexibility, and explain why heterogeneity can sometimes reduce correlated failure without being intrinsically desirable.

The argument is intentionally narrower than a general theory of resilience. Classical resilience theory distinguished persistence through disturbance from rapid return to a fixed equilibrium (Holling 1973). Adaptive governance research subsequently emphasized learning, networks, multiple forms of knowledge, and reorganization during periods of abrupt change (Folke et al. 2005). The present framework uses these literatures as conceptual resources while focusing on a more specific problem: what kinds of response capacity matter when susceptibility is elevated and the political meaning of a disturbance is not settled in advance.

Table 1 makes one point explicit: suppression is a member of the response repertoire rather than its default organizing principle. The table also prevents the opposite simplification. Observation, restraint, experimentation, and amplification are not inherently superior. Their appropriateness depends on the receiving configuration, the expected irreversibility of consequences, the distribution of risks and benefits, and the legitimacy of the actors selecting the response.

9.1 Response repertoire as a property of the system

Let η_t denote a perturbation encountered at time t. Define the feasible response repertoire as

$$\mathcal{A}_t(\eta_t) = \left\{ a_t^{(1)}, a_t^{(2)}, \ldots, a_t^{(m)} \right\}, \tag{43}$$

where each $a_t^{(k)}$ denotes a materially available response mode or combination of modes. Feasibility is political as well as technical. A response may be dynamically conceivable yet unavailable because institutions lack the relevant authority, resources, coordination, knowledge, trust, or legal permission. Conversely, a formally available instrument may be unusable in practice because its governance footprint would be unacceptable.

The size of $\mathcal{A}_t$ is therefore only a crude indicator. A system with many nominal policy instruments can possess a poor effective repertoire when all instruments operate through the same vulnerable channel or generate highly correlated side effects. A smaller repertoire can be more adaptive if its members are genuinely distinct, reversible, and supported by heterogeneous information sources. The relevant object is thus the structure of available responses rather than a count of policy tools.

Cybernetic work on requisite variety provides a useful point of comparison. Ashby's formulation links the capacity of a regulator to the variety of possible disturbances, while Conant and Ashby later formalized conditions under which successful regulation requires a model of the system being regulated (Ashby 1956; Conant and Ashby 1970). The political framework here retains only a restrained lesson: a response system with extremely narrow variety can be overwhelmed by disturbances that require qualitatively different forms of reception. The stronger cybernetic objective of maintaining specified essential variables cannot be imported without modification, because political systems contain contested ends, reflexive actors, unequal burdens, and uncertain normative boundaries.

Definition 9.1 (Response repertoire)**.** A *response repertoire* is the historically and institutionally available set of distinct ways through which a political configuration can receive, transform, delay, redirect,

Table 1: Provisional response repertoire for near-critical political settings. The modes are analytical possibilities rather than a universal sequence, and several can coexist within one response.

Response mode	Primary dynamical role	Potential epistemic value	Characteristic risk
Observation	Preserve information before major intervention	Reveal susceptibility, coupling, and endogenous adaptation	Delay can permit accumulating harm or irreversible transition
Absorption / buffering	Reduce transmission of a perturbation	Test whether disturbance remains locally containable	Buffering can hide stress or protect an unjust configuration
Decoupling	Interrupt propagation across subsystems	Identify dependence on network pathways	Fragmentation can damage coordination and shift burdens elsewhere
Redirection	Alter the pathway through which effects travel	Reveal alternative channels and response structures	Displaced effects can become less visible rather than less harmful
Delay	Slow escalation or defer commitment	Create time for observation, consultation, and model revision	Delay can become strategic avoidance or allow lock-in to deepen
Experimentation	Introduce bounded, observable, revisable change	Generate information about response functions	Even small experiments can cross thresholds under high susceptibility
Selective amplification	Strengthen a process already emerging	Test or support potentially generative dynamics	Amplification can create cascades, exclusion, or premature trajectory closure
Restraint / permission	Refrain from forcing a specific trajectory	Observe self-organization and preserve option space	Inaction can be interpreted, distributive, and dynamically consequential
Strong suppression	Rapidly reduce a perturbation judged intolerable	Protect boundary conditions under urgent threat	Suppression can destroy useful variation, information, or political agency

absorb, amplify, or refrain from amplifying a perturbation. Its adequacy depends on the relation between perturbation variety, response diversity, switching capacity, and the normative admissibility of the available responses.

This definition shifts adaptive capacity away from the image of a powerful central controller. The capacity can be distributed across organizations, jurisdictions, professional communities, informal networks, and affected populations. Adaptive-governance research similarly emphasizes connections across organizational levels and the contribution of multiple knowledge systems and bridging organizations during periods of change (Folke et al. 2005). The present argument adds that such distribution matters especially when one response channel can become amplified near a critical region.

9.2 Response diversity and correlated failure

Near criticality, synchronized response can become a source of vulnerability. Suppose n subsystems encounter a common perturbation and all implement nearly identical responses. Let

$$\rho_{ij}^{R} = \mathrm{Corr}(R_i(\eta_t), R_j(\eta_t)) \tag{44}$$

represent the correlation between subsystem responses in an illustrative statistical description. High average ρ_{ij}^{R} can increase the probability that a mistaken classification, shared model error, or common incentive propagates through the entire system at once. Heterogeneous responses can reduce this form of correlated failure by allowing some units to absorb, some to observe, some to experiment, and some to preserve continuity.

The argument should not be converted into a universal preference for fragmentation. Heterogeneity can also obstruct coordination, reproduce inequality, create jurisdictional arbitrage, or delay responses to urgent collective threats. Polycentric-governance scholarship provides an important institutional comparison because it treats multiple formally independent centers of decision making as potentially capable of mutual adjustment and learning rather than assuming that complex problems require one center (Ostrom 2010). The value of such plurality remains conditional on communication, accountability, and the structure of the problem.

Proposition 9.2 (Correlated-response vulnerability)**.** *When a near-critical political system is exposed to common uncertainty, strong homogeneity in observation models and response channels can increase exposure to correlated error and synchronized amplification. Response heterogeneity can reduce that exposure when distinct units remain sufficiently connected to share information and when the heterogeneous responses do not generate larger coordination or justice failures.*

The proposition frames heterogeneity as a portfolio property rather than an intrinsic virtue. Diversification has value because errors, delays, and responses need not be perfectly correlated. The same logic implies that apparent institutional plurality offers little protection when different units share the same data pipeline, model assumptions, funding incentives, or political pressure. Near-critical analysis therefore asks about substantive response independence rather than organizational labels alone.

9.3 Adaptive capacity and switching capacity

A repertoire matters only if a system can move among its members. Let $\kappa_t(a_i, a_j)$ denote the practical cost of switching from response mode a_i to a_j. High switching costs can convert an apparently rich repertoire into a rigid response pattern. Lock-in can arise through law, contracts, organizational routines, infrastructure, professional doctrine, sunk costs, or political commitments. The path-dependence mechanisms discussed in Section 5 therefore apply to response systems as well as to broader political trajectories.

A provisional representation of adaptive capacity is

$$\mathcal{C}_t^{\mathrm{A}} = \Psi(\mathcal{A}_t, \mathcal{I}_t, \mathcal{K}_t, \mathcal{L}_t, \mathcal{V}_t, \mathcal{J}_t)\,, \tag{45}$$

where $\mathcal{I}_t$ denotes informational diversity, $\mathcal{K}_t$ switching capacity, $\mathcal{L}_t$ learning capacity, $\mathcal{V}_t$ reversibility or option-preserving features, and $\mathcal{J}_t$ the normative and institutional conditions that make responses legitimate and distributively defensible. The function Ψ is deliberately unspecified. Equation (45) is a conceptual inventory rather than a proposed estimator.

Remark 9.3 (Adaptive capacity)**.** Adaptive capacity near criticality concerns the ability to change how a system responds while consequences are unfolding. It differs from possessing many instruments, from returning rapidly to a previous state, and from maximizing intervention speed. A response that cannot be revised after new information arrives can have low adaptive value even when its immediate effect is large.

This distinction echoes Holling's separation between resilience and a narrow notion of stability (Holling 1973). In the present framework, rapid return to a prior state can be undesirable when the prior state is itself unjust, brittle, or locked into an unfavorable attractor. Adaptive capacity is therefore tied to continued responsiveness rather than restoration alone.

9.4 Response sequencing and reversible escalation

Response modes need not be selected once. They can be sequenced. A system might begin with observation and buffering, move to bounded experimentation, and escalate toward stronger suppression only if evidence of catastrophic propagation accumulates. Another system might require immediate boundary protection while preserving room for later redirection and institutional learning. The sequence is part of the response because earlier steps modify the information and option space available to later ones.

Let a response path be

$$\Pi_t = \left(a_t^{(1)}, a_{t+1}^{(2)}, \ldots, a_{t+h}^{(q)}\right). \tag{46}$$

A near-critical assessment of Π_t should consider at least three features: its capacity to slow destructive amplification, its production of information, and its reversibility if the working interpretation proves mistaken. This does not imply a universal gradualism. Some hazards can justify rapid and forceful boundary protection. The point is that escalation speed itself should be analyzed as part of the governance footprint.

Working Claim 9.4 (Reversible escalation)**.** Where immediate catastrophic harm is not already dominant, a response sequence that preserves observation and revision between stages can possess greater adaptive value than a one-step commitment of comparable intended effect. The claim is conditional on the intermediate stages remaining dynamically and ethically tolerable.

The qualification is essential. Deliberate gradualism can externalize risk onto actors who already face severe harm. Reversibility for the governing institution can coexist with irreversibility for affected persons. The later discussion of transition justice will therefore return to the distribution of transition burdens and to the question of whose options are preserved by a supposedly flexible response.

9.5 Learning through perturbation

The repertoire framework also changes the epistemic status of ordinary fluctuations. A disturbance can generate information about which response modes remain viable. If a buffer absorbs the perturbation without secondary effects, the system learns something different than if buffering shifts stress into a connected subsystem. If local experimentation produces divergent outcomes, that divergence can reveal hidden heterogeneity. Perturbation becomes part of a learning relation rather than merely an object to suppress.

This interpretation is compatible with the embeddedness developed in Section 8. Learning itself alters the field. Actors can imitate an apparently successful response, resist it, or strategically exploit new rules. Adaptive capacity therefore includes the ability to update without converting every provisional lesson into synchronized system-wide action.

Definition 9.5 (Response learning)**.** *Response learning* is the revision of the feasible repertoire, switching rules, or interpretation of perturbations through observed consequences of prior responses. It is distinct from simple policy feedback because the object of learning includes the response architecture itself: which modes remain available, how they interact, and when they become dangerous near criticality.

The distinction is relevant to complex adaptive governance. Duit and Galaz emphasize the difficulty of governing systems characterized by nonlinear change, uncertainty, and self-organization (Duit and Galaz 2008), while May's framework treats governance arrangements, organizations, and people as interacting components of complex adaptive systems (May 2022). The present concept of response learning adds a specific near-critical question: can institutions revise their own mode of response before one mode becomes self-reinforcing and closes alternatives?

9.6 Response capacity and non-coercive orientation

The response repertoire offers a bridge to the paper's later discussion of *wuwei*. A system with only a coercive suppression mode has limited adaptive variety. Yet a system that refuses strong intervention under all conditions is equally narrow. Non-coercive orientation therefore cannot be identified with passivity. It concerns the capacity to receive perturbation through modes whose force is proportionate to the evolving relation and whose use does not prematurely assume a fully specified destination.

The water metaphor developed in the introductory discussion becomes more precise at this point. Water encountering heterogeneous terrain displays a range of possible responses: diversion, acceleration, pooling, dissipation, and continued flow. Political systems should not be naturalized through this image, and human harms cannot be treated as aesthetically interesting turbulence. The analogy is useful only for loosening the assumption that an obstacle has one proper response, namely removal. A response repertoire allows the system to remain active without requiring every disturbance to be framed immediately as an enemy of order.

Proposition 9.6 (Repertoire principle)**.** *Near criticality, adaptive political capacity is better represented by the availability, differentiation, revisability, and normative admissibility of multiple response modes than by the maximum strength of a single steering instrument. This principle does not determine which response is justified in a particular case; it identifies a structural condition under which responsive judgment remains possible.*

The proposition prepares the next section. Once a repertoire of responses is recognized, the political question becomes how action can remain responsive without turning every available capacity into an imperative to steer. The next section examines *wuwei* and non-coercive response as a philosophical resource for thinking about action that remains situated within evolving conditions rather than beginning from a fully specified terminal state.

10 Wuwei and Non-Coercive Response

This section introduces *wuwei* as a comparative philosophical resource for clarifying a problem generated by the preceding sections. Near criticality, a political system may exhibit elevated susceptibility, uncertain trajectory closure, and a wide gap between local epistemic leverage and legitimate action. The availability of a response repertoire therefore leaves a further question: how can political action remain capable of intervention while avoiding the assumption that every perturbation should be converted into an object of steering? The section develops three ideas. First, *wuwei* is interpreted as an orientation toward action that reduces assertive forcing and attends to conditions already present in the field. Second, the paper introduces *non-coercive response* as a distinct analytical category situated between passive inaction and destination-driven steering. Third, the section specifies ethical limits to this orientation, especially where waiting, yielding, or minimal intervention transfers severe risk to actors with less capacity to absorb harm. The purpose is conceptual clarification rather than direct translation of classical Chinese philosophy into a modern policy doctrine.

10.1 Interpretive scope of wuwei

The term *wuwei* is commonly rendered as "nonaction," "effortless action," "non-assertive action," or "non-coercive action." Contemporary scholarship emphasizes that the term does not denote simple inactivity. In the *Laozi*, it is connected with *ziran*, naturalness or self-so-ness, and with a mode of action less driven by assertive desire and self-conscious striving (Chan 2025). Slingerland's comparative study of early Chinese thought likewise treats *wuwei* as a family of ideals concerning effortless or unselfconscious action rather than literal behavioral absence (Slingerland 2003). More recently, Valmisa has interpreted *wuwei* through relational and distributed agency, emphasizing action that emerges through assemblies and multiple determinations (Valmisa 2025).

These interpretations are not interchangeable, and the present paper does not attempt to adjudicate among them. They nevertheless identify an overlapping philosophical space useful for the current problem: action can remain effective without beginning from a maximally explicit command over the environment. Such a possibility matters near criticality because the demand for control can become strongest at precisely the moment when the consequences of control are least reliably bounded.

Remark 10.1 (Comparative use)**.** The paper uses *wuwei* to illuminate a structural tension between forceful steering and responsive participation. It does not claim that ancient Chinese political philosophy anticipated contemporary critical-transition theory, nor that a modern institution can operationalize *wuwei* through a fixed set of procedures. The comparison is heuristic and philosophical.

The distinction also prevents an easy assimilation of *wuwei* to institutional passivity. A government that withholds action while preventable harm accumulates can still exercise power through delay, agenda control, or unequal exposure to risk. Conversely, an institution can act extensively while reducing coercive closure if its actions preserve multiple pathways, remain revisable, and respond to information generated by the evolving field. The relevant axis therefore concerns the relation between action and forcing, not the visible quantity of activity alone.

10.2 Three modes of political response

The preceding distinction can be formalized through three ideal types. Let η_t denote a perturbation, x_t the current configuration, and u_t an intentional political response. The ideal types differ in the way u_t is related to the future configuration.

Table 2 summarizes the three ideal types and their characteristic risks. Passive inaction leaves the response architecture largely unchanged. Coercive steering evaluates perturbations mainly through their

Table 2: Three ideal-typical modes of response near criticality. The categories are analytical and may overlap in practice.

Mode	Orientation	Relation to perturbation	Characteristic risk
Passive inaction	Minimal intentional response	Perturbation proceeds with little deliberate adjustment	Existing asymmetries and avoidable harms can deepen during delay
Coercive steering	Movement toward a specified target or regime	Perturbation is suppressed, amplified, or redirected according to a prior destination	Model error and normative error can be amplified into trajectory closure
Non-coercive response	Situated participation with limited closure pressure	Perturbation is observed, received, buffered, redirected, or joined according to evolving relations	Ambiguity can become indecision, and restraint can externalize risk

relation to a specified end state. Non-coercive response occupies a different orientation. It treats perturbations as events whose meaning can change as they interact with susceptibility, coupling, history, and other actors. Intervention remains possible, including strong intervention where boundary protection becomes necessary, while the response begins with a lower commitment to a predetermined terminal configuration.

Definition 10.2 (Non-coercive response)**.** A *non-coercive response* is an intentional participation in an evolving political field that seeks to influence immediate conditions while limiting premature trajectory closure. Its characteristic features are attentiveness to emergent response, proportionality of force, revisability where feasible, and continued recognition that the political meaning of a perturbation can remain partly unsettled during action.

The definition contains an important asymmetry. Non-coercive response does not promise low material force in every case. A rapidly propagating hazard may justify forceful containment. The adjective *non-coercive* refers to the orientation toward closure: the response avoids converting available leverage into a general entitlement to determine the system's long-run attractor. This interpretation keeps emergency boundary protection conceptually separate from destination-driven social engineering.

10.3 Forcing and trajectory closure

The idea can be represented by distinguishing intervention magnitude from *closure pressure*. Let $\|u_t\|$ denote a rough measure of intervention magnitude and let $\kappa(u_t, x_t)$ denote the degree to which the intervention contracts the set of practically reachable future trajectories. The two need not move together. A small intervention near a sensitive boundary can generate large closure pressure,

$$\|u_t\| \text{ small}, \qquad \kappa(u_t, x_t) \text{ large}, \tag{47}$$

while a larger protective intervention can sometimes preserve a wider future option space by preventing an irreversible cascade.

This distinction changes the interpretation of restraint. Restraint concerns more than reducing $\|u_t\|$. It can also mean limiting κ: avoiding choices that exploit high susceptibility to close alternatives whose ethical or political status remains unsettled. Near criticality, a low-magnitude action can therefore be highly assertive in historical consequence.

Working Claim 10.3 (Closure-sensitive restraint)**.** When uncertainty about long-run consequences and normative standing remains high, restraint should be evaluated partly by the intervention's expected con-

tribution to trajectory closure rather than by intervention magnitude alone. This claim remains defeasible where delay itself creates severe or irreversible harm.

The claim connects directly to Section 7. A well-estimated local direction can identify a path with large dynamical leverage. The same information can raise reasons for caution because leverage increases the capacity to close trajectories. *Wuwei*, interpreted here as reduced assertive forcing, supplies a philosophical vocabulary for the possibility that knowing how to push effectively need not create an obligation to push.

10.4 Water and responsive form

Water imagery in the *Laozi* is associated with yieldingness, low position, and a form of strength that does not depend on hardness or direct contention (Chan 2025). The image is useful for the present framework because it directs attention from obstacle elimination to response formation. When a flow encounters heterogeneous terrain, its resulting pattern depends on the relation among velocity, geometry, resistance, neighboring flows, and prior configuration. The obstacle enters the formation of the flow without becoming the sole determinant of its direction.

The political analogy can be stated cautiously. Let G_t represent a receiving configuration and $R(\eta_t, G_t)$ the response generated when perturbation η_t enters it. A response-centered orientation asks how the configuration receives the perturbation before assigning a fixed political meaning to the perturbation itself. The relevant object remains the pair

$$(\eta_t, R(\eta_t, G_t)), \tag{48}$$

which was introduced in Section 4. The water metaphor adds a practical intuition: response can involve division, temporary storage, dissipation, redirection, or continued movement around resistance. Removal is one possible relation among several.

Proposition 10.4 (Responsive-form principle)**.** *Near criticality, the political significance of a perturbation is partly constituted through the form of response that emerges around it. Institutions therefore benefit from capacities for receiving and differentiating perturbations before committing every disturbance to the same response mode.*

The proposition does not aestheticize suffering. Political perturbations can include violence, deprivation, exclusion, or rapidly compounding material loss. The river image has no authority to diminish those harms. Its analytical value is narrower: it weakens the assumption that effective response always requires direct opposition to the initiating disturbance.

10.5 Responsive participation

The concept of *responsive participation* gives the paper a vocabulary that can travel beyond the Daoist comparison. Governance was described in Section 8 as embedded within the same field it seeks to influence. Responsive participation accepts that embeddedness as a starting condition. The political actor observes, acts, receives the consequences of action, and revises its position while recognizing that each response can alter the subsequent field of possibilities.

A simplified response cycle can be represented by the ordered configuration

$$\mathcal{C}_t = (\eta_t, O_t, u_t, R_t, O_{t+1}), \tag{49}$$

where O_t denotes provisional observation, u_t an intentional response, and R_t the distributed response generated through the field. The ordering marks stages of attention and response without assuming a

deterministic causal chain. The next observation is taken in a system partly transformed by the previous response.

This cycle differs from an optimization architecture in which a stable model maps observed states to a target-maximizing action. Responsive participation places greater weight on intermediate observation, field-generated information, and the possibility that the relevant value judgment changes as previously latent consequences become manifest. The orientation is especially important when the available attractors carry heterogeneous distributions of benefit, risk, and agency.

Working Claim 10.5 (Provisionality of destination)**.** Where several politically viable trajectories remain open and their distributive or ethical consequences are incompletely manifest, a response can possess political value through preserving the possibility of later judgment. The claim supports provisionality of destination under specified uncertainty; it does not establish openness as an overriding value.

This formulation creates a bridge to the later sections on justice and revisability. A trajectory can remain open while current harms are distributed very unequally. Preserving future options for a political system can therefore coexist with closing options for particular persons. The aggregate vocabulary of openness must be disaggregated before it can acquire normative force.

10.6 Asymmetric vulnerability and the limits of restraint

The strongest objection to a non-coercive orientation concerns asymmetric vulnerability. Suppose actor A can safely wait while actor B bears accumulating loss during the same period. An institutional decision to "observe further" then has a distributive structure. Delay can preserve options for the institution while consuming the options of the more exposed actor. A language of patience, yielding, or spontaneous adjustment can conceal this asymmetry if it is detached from the distribution of risk.

For this reason, the paper treats *wuwei* as a constraint on assertive forcing and not as a general presumption for minimal action. The relevant assessment must include at least the rate at which harms accumulate, the distribution of exposure, the reversibility of delay, and the capacity of vulnerable actors to participate in subsequent revision.

Research Obligation 10.6 (Asymmetry test)**.** Any application of non-coercive response should identify whose risks are reduced by restraint, whose risks are extended by it, and whether the preservation of system-level option space is purchased through the contraction of option space for less protected actors. This test belongs to the normative analysis and cannot be inferred from dynamical susceptibility alone.

The asymmetry test also blocks a romantic reading of self-organization. Self-organizing processes can reproduce hierarchy, exclusion, or accumulated advantage. Decentralized emergence supplies information about how order forms; it does not supply a sufficient criterion of justice. The distinction prepares Sections 12 and 13, where the paper turns explicitly to the distribution and moral weight of transition consequences.

10.7 Non-coercive orientation and the meaning of governance

The analysis now permits a narrower use of the term *governance*. Near criticality, governance need not denote the successful imposition of a terminal state. It can denote organized capacities for sensing, receiving, buffering, responding, revising, and protecting boundaries within an evolving field. This usage remains normative because decisions about protection, admissible harm, and preserved options require political judgment. Its difference lies in the location of normativity: the framework initially constrains modes of participation and closure while postponing a complete specification of the future order.

Proposition 10.7 (Responsive-participation principle)**.** *Near criticality, a governance orientation can remain politically active while limiting premature trajectory closure by combining differentiated response,*

continued observation, proportionality, reversibility where feasible, and explicit attention to asymmetric vulnerability. These properties define a mode of participation; they do not determine the justice of a particular outcome.

This principle offers a provisional answer to the question that motivated the section. *Wuwei* contributes a philosophical language for effective action with reduced assertive forcing, while the response framework translates that intuition into analyzable dimensions without treating ancient concepts as ready-made policy instruments. The next section develops the adjacent concept of *shi*: the propensity that arises from configuration. That move shifts attention from the actor's chosen force toward the directional tendencies already forming within the political field.

11 Configuration, Propensity, and Shi

This section develops a configurational vocabulary for the directional tendencies that can arise before any political actor selects a target state. Sections 10 and 9 shifted attention from forceful steering toward differentiated response. The present section asks what such response is responding *with*: the patterned relations, accumulated histories, institutional couplings, material constraints, and actor expectations that make some movements comparatively easy to amplify and others comparatively difficult to sustain. The classical Chinese term *shi* provides a comparative resource for this problem because several of its historical and modern interpretations connect efficacy to the configuration of a situation and to situational disposition beyond isolated force. The section proceeds in five steps. It first limits the comparative use of *shi*; it then defines a configuration-dependent propensity profile; distinguishes propensity from intervention force and from prediction; examines competing propensities near criticality; and concludes by separating configurational efficacy from normative legitimacy. The objective is a vocabulary for tendencies already forming in a political field, together with safeguards against turning those tendencies into a new deterministic or technocratic doctrine.

11.1 Interpretive scope of shi

Shi has carried multiple meanings across Chinese intellectual history, including position, circumstance, power, strategic advantage, momentum, and tendency. Its military uses in texts associated with the *Sunzi* are closely connected to strategic configuration and to advantages generated by the relation among force, terrain, timing, and deployment. Raphals documents this configurational vocabulary in her comparison of the *Sunzi* and *Xunzi*, including the language of strategic advantage associated with *shi* (Raphals 2016). Jullien's influential comparative reconstruction develops a broader reading of *shi* as the "propensity of things," emphasizing efficacy that arises from the disposition of a situation (Jullien 1995). More recent work has experimented with *shi* as a concept for social-scientific description and organizational process (Law and Lin 2018; Yue et al. 2023).

These uses form a conceptual family across differing historical and interpretive contexts. The present paper therefore uses *shi* in a narrow comparative sense: *configuration-dependent propensity*. The term identifies the possibility that a field can contain directional tendencies generated by its current relations, histories, and constraints. This usage preserves the configurational insight while leaving the historical interpretation of classical texts to specialist scholarship.

Remark 11.1 (Comparative discipline)**.** The paper treats *shi* as a source of conceptual contrast and analytical vocabulary. It avoids treating "Chinese thought" as a uniform tradition and avoids using *shi* as evidence that political systems possess a measurable quantity equivalent to a physical force. Its role is heuristic: to direct attention toward efficacy that can arise from configuration.

This caution matters because the concept could otherwise be absorbed too easily into the control language criticized earlier. If *shi* were translated as an invisible force that a skilled governor merely detects and exploits, the paper would reproduce a more subtle version of destination-driven steering. A configurational reading asks how tendencies are distributed across the field, how they change as actors respond, and how several tendencies can coexist while leaving multiple futures open.

11.2 Configuration-dependent propensity

Let the political configuration at time t be represented schematically as

$$\mathcal{C}_t = (x_t, G_t, H_t, \Theta_t), \tag{50}$$

where x_t denotes the currently observed state variables, G_t the relevant coupling or interaction structure, H_t the historically accumulated path, and Θ_t slower-moving institutional, material, or normative condi-

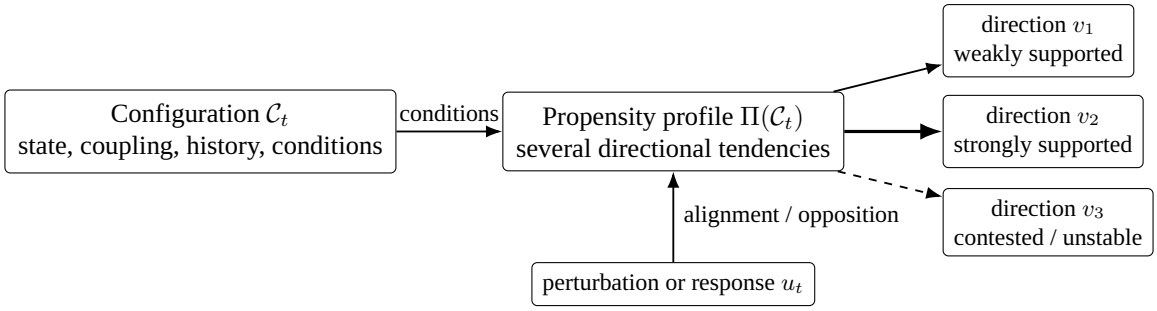

Figure 8: Configuration-dependent propensity. The same perturbation can acquire different effects because the receiving configuration supports several directional tendencies with different strengths and uncertainties. Line weights are illustrative only.

tions. The components are analytical placeholders. Their purpose is to prevent propensity from being attributed to the visible state alone.

A *propensity profile* can then be represented as

$$\Pi(\mathcal{C}_t;\tau) = \{(v, \rho_t(v;\tau)) : v \in \mathcal{V}_t\}, \tag{51}$$

where $\mathcal{V}_t$ is a set of politically meaningful directions over a horizon τ, and $\rho_t(v;\tau)$ is a provisional measure of the relative tendency for change in direction v. The symbol ρ admits probabilistic and non-probabilistic interpretations. In empirical applications it could be estimated probabilistically, qualitatively, or through multiple indicators. The formal expression serves to make one point explicit: a configuration can support several directional tendencies with different strengths and uncertainties.

Definition 11.2 (Configuration-dependent propensity)**.** A *configuration-dependent propensity* is a directional tendency whose relative ease of amplification or persistence arises from the current relation among system state, coupling structure, historical path, and slower-moving conditions. It describes an orientation of possible change while leaving the eventual outcome open.

Figure 8 summarizes the distinction. The configuration generates a profile of relative tendencies; an intervention then enters that already structured field. The analytical priority therefore shifts from the intervention as an isolated cause toward the relation between the intervention and the configurational tendencies it encounters.

This formulation also clarifies the relation to Sections 4 and 6. Susceptibility concerns how strongly the system responds in specified directions. Propensity concerns the directional tendencies already supported by the broader configuration. The two can interact while remaining analytically distinct. A system can display high susceptibility while its directional future remains contested, and a strong historical tendency can persist while short-horizon susceptibility remains modest.

11.3 Propensity, force, and configuration leverage

A force-centered model asks how large an intervention must be to move the system. A propensity-centered model additionally asks whether the intervention works with, across, or against tendencies already generated by the field. Let u_t denote an intentional intervention and let D_τ measure the divergence between a trajectory following that intervention and a comparison trajectory. An illustrative measure of configuration leverage is

$$\Lambda_t(u_t;\tau) = \frac{D_\tau\left(x^{(u)}_{t+\tau}, x^{(0)}_{t+\tau}\right)}{\|u_t\|}, \tag{52}$$

where the ratio is meaningful only when the numerator and intervention norm have an interpretable empirical specification. The formal sketch shows why intervention magnitude and trajectory consequence can separate. Alignment with a strongly supported tendency can produce high leverage; opposition to a strongly supported tendency can require sustained effort; and near a basin boundary a very small action can yield high leverage even when the underlying propensity profile remains uncertain.

Proposition 11.3 (Configurational leverage)**.** *The consequence of an intervention near criticality depends partly on its relation to configuration-dependent propensities. Intervention magnitude alone therefore supplies an incomplete account of political efficacy. This proposition concerns dynamical leverage and leaves normative evaluation open.*

The last sentence is essential. A political tendency can be efficient to join and still be ethically troubling. A self-reinforcing hierarchy, exclusionary institution, speculative cascade, or coercive coalition can develop strong configurational momentum. The ease of moving with that momentum supplies information about efficacy; it supplies only limited grounds for normative justification.

The distinction also complicates the language of "creating" propensity. Political actors can alter networks, expectations, institutional rules, resource flows, and public narratives in ways that change $\mathcal{C}_t$ and hence its propensity profile. Such configurational intervention can sometimes achieve large downstream effects with limited direct force. Modern organizational readings of *shi* explicitly distinguish creating favorable momentum from leveraging existing momentum (Yue et al. 2023). In the present framework, configuration-shaping remains a form of intervention and therefore remains subject to the same questions of legitimacy, distribution, and trajectory closure developed throughout the paper.

Working Claim 11.4 (Configuration-shaping equivalence)**.** Altering the configuration from which future tendencies emerge can carry political consequences comparable to direct state intervention. Lower visible force therefore provides limited evidence about lower normative significance.

This claim prevents a soft form of manipulation from escaping scrutiny. A political actor who changes information architecture, participation rules, or resource access may influence subsequent self-organization through indirect mechanisms, including changes that occur before any visible command is issued. Near criticality, such indirect changes can be especially consequential because the field is already sensitive to shifts in coupling and expectation.

11.4 Competing propensities near criticality

Near criticality, the propensity profile can become plural and unstable. A configuration may support several incipient directions whose relative strength changes as perturbations arrive. This possibility is especially important for the paper's treatment of random noise. A fluctuation can temporarily amplify one incipient direction, alter expectations, and thereby change the subsequent propensity profile. The field then exhibits a recursive relation,

$$\Pi(\mathcal{C}_t) \longmapsto R(\eta_t, \mathcal{C}_t) \longmapsto \mathcal{C}_{t+1} \longmapsto \Pi(\mathcal{C}_{t+1}), \tag{53}$$

where each response can modify the configuration from which later tendencies arise. Equation (53) is a conceptual ordering that keeps transition contingent.

This recursive structure makes the phrase "follow the tendency" too simple for near-critical politics. First, several tendencies can coexist. Second, the act of publicly identifying a tendency can change actor expectations and hence the tendency itself. Third, short-horizon momentum can point toward a trajectory that later becomes difficult to reverse. Fourth, distributed actors can perceive different propensities from different positions in the system. The political field therefore resembles a changing propensity landscape more than a single current with one obvious direction.

Definition 11.5 (Propensity plurality)**.** *Propensity plurality* describes a configuration in which several directional tendencies remain materially or politically plausible and their relative support is sensitive to perturbation, interpretation, or endogenous response. The concept marks a limit to strategies that infer one dominant future from current momentum.

Propensity plurality also clarifies the role of noise developed in Section 3. Noise can reveal which tendencies are easily activated, can redistribute relative support among them, and can become incorporated into new feedback structures. The appropriate response therefore depends on more than whether a disturbance aligns with the strongest presently observed tendency. The response must also consider how that alignment changes the field and which alternatives it may close.

11.5 Shi and non-coercive response

The configurational interpretation of *shi* helps specify the water intuition introduced in Section 10. Water acquires a path through its relation to gradient, terrain, obstacles, volume, and prior channels. Political response can similarly attend to relations that make some adjustments easier to sustain than others. This analogy supports an orientation toward working with the receiving field, while the differences between hydrodynamic and political systems remain decisive. Political actors interpret their situation, contest descriptions of it, organize against prevailing tendencies, and make claims of justice that are specific to human political relations.

Law and Lin's experimental use of *shi* in social science is valuable here because it emphasizes situated flows and counterflows and treats inquiry itself as participating in the context it describes (Law and Lin 2018). That emphasis fits the embedded-governance argument of Section 8. The practical implication is modest: political response can include the diagnosis of configurational propensities before selecting a response mode, while preserving awareness that diagnosis can itself enter the field.

A provisional response sequence can therefore be written as

$$\mathcal{C}_t \longrightarrow \widehat{\Pi}_t \longrightarrow \mathcal{R}_t \longrightarrow \text{response selection} \longrightarrow \mathcal{C}_{t+1}, \tag{54}$$

where $\widehat{\Pi}_t$ denotes a fallible inference about propensity and $\mathcal{R}_t$ the available response repertoire. The sequence gives observation and configurational diagnosis a place before intervention, while its recursive endpoint prevents the diagnosis from being treated as permanent.

Working Claim 11.6 (Propensity-sensitive response)**.** A non-coercive orientation near criticality can use information about configuration-dependent propensity to select among response modes while preserving uncertainty about the final trajectory. Propensity information can support timing, buffering, redirection, or restraint while leaving the terminal attractor normatively unsettled.

11.6 Efficacy and normative standing

The concept of *shi* creates a final danger that must be resolved before the paper turns to justice. Political actors can easily transform a descriptive claim about momentum into a normative claim about legitimacy. A trajectory can appear inevitable, efficient, historically mature, or aligned with the "direction of events," and those descriptions can be used to silence actors whose interests are harmed by the emerging configuration. The framework therefore separates three judgments:

$$\text{propensity judgment} \quad ; \quad \text{efficacy judgment} \quad ; \quad \text{normative judgment}. \tag{55}$$

A propensity judgment concerns what tendencies the configuration appears to support. An efficacy judgment concerns how a response may interact with those tendencies. A normative judgment concerns whether participating in, redirecting, or resisting the tendency is justified.

Proposition 11.7 (Normative independence of propensity)**.** *A strong configuration-dependent propensity supplies evidence about political dynamics and possible leverage. Its strength alone supplies insufficient grounds for justice, legitimacy, or ethical permissibility. Near criticality, the normative importance of this distinction increases because alignment with an emerging tendency can contribute to persistent trajectory closure.*

This proposition returns the paper to the issue that first separated its argument from control theory. The capacity to identify a powerful tendency can increase intervention capability while simultaneously increasing responsibility for the futures that such intervention helps close. A *shi*-sensitive framework therefore has two complementary uses. It directs attention toward configuration beyond isolated force, and it reveals why efficacy cannot serve as its own normative criterion.

Research Obligation 11.8 (Propensity audit)**.** Before configurational momentum is used to justify political action, analysis should distinguish the actors who benefit from the tendency, the actors whose options contract as it strengthens, the mechanisms that reproduce the tendency, and the degree to which alternative trajectories remain practically reachable. These questions provide the bridge from configurational dynamics to the justice of political transitions.

The next section develops that bridge directly. Once near-critical response is understood through competing propensities, high leverage, and possible trajectory closure, the central political question concerns the distribution of benefits, burdens, agency, and future options across the transition. The analysis therefore moves from the efficacy of configuration to the justice of political becoming.

12 Justice of Political Transitions

This section supplies the normative bridge required by the preceding dynamical analysis. Sections 3–11 developed concepts for susceptibility, response fields, trajectory closure, embedded intervention, response repertoires, non-coercive response, and configuration-dependent propensity. Those concepts can identify conditions in which political action has unusual leverage. They leave open the political question of whether that leverage should be exercised, by whom, through which procedures, and with what distribution of consequences. The structural role of this section is therefore to separate dynamical reachability from normative standing and to identify dimensions along which a near-critical transition can be evaluated as a political transition rather than only as a change of state.

The discussion proceeds in four steps. First, it distinguishes the justice of a resulting configuration from the justice of the path through which that configuration is produced. Second, it introduces distributional incidence, agency, legitimacy, option-space effects, reversibility, and temporal incidence as analytically distinct dimensions. Third, it examines how elevated susceptibility can amplify normative consequences even when the initiating intervention is small. Fourth, it develops a provisional transition-justice profile that preserves disagreement among competing theories of justice rather than reducing them to a single optimization criterion.

The familiar literature on distributive justice already treats political and institutional arrangements as consequential for the allocation of benefits and burdens, while allowing substantial disagreement about the appropriate principles of distribution (Lamont and Favor 2017). Political legitimacy raises a related but distinct set of questions concerning the justification of authority, coercion, procedure, participation, and substantive outcomes (Peter 2023). Near-criticality adds a dynamical complication to these debates: an intervention can affect the distribution of present goods while also changing which future institutional and political configurations remain practically reachable.

12.1 Transition states and transition paths

Let x_t denote a political configuration in a near-critical region, and let u_t denote an intentional intervention whose interaction with the receiving field contributes to a transition toward configuration x_{t+1}. A purely state-centered evaluation might compare the properties of x_t and x_{t+1}. A transition-centered evaluation also examines the path

$$\gamma_{t:t+k} = \{x_t, u_t, R_t, x_{t+1}, \ldots, x_{t+k}\}, \tag{56}$$

where R_t represents the distributed responses through which the intervention is interpreted, resisted, amplified, redirected, or incorporated into the political field. The path records politically relevant features that can be lost when analysis moves directly from one state description to another.

A transition can improve some aggregate or institutional property while producing burdens through exclusion, coercion, unequal risk imposition, or asymmetric contraction of future options. Conversely, a procedurally inclusive transition can still generate a resulting configuration with serious substantive injustice. Political legitimacy and justice are therefore related without being analytically identical; contemporary legitimacy theory contains both procedural and substantive approaches, and it commonly distinguishes legitimate authority from mere effectiveness (Peter 2023). The near-critical framework adds a temporal reason to preserve that distinction: when transition effects become persistent, the mode of arrival can shape who is able to contest, revise, or exit the resulting order.

Definition 12.1 (Transition justice)**.** *Transition justice* denotes the normative evaluation of a political change as a temporally extended process. It includes the distribution of benefits and burdens, the standing and participation of affected actors, the means through which the transition is produced, the contraction or expansion of future option spaces, the reversibility of resulting commitments, and the allocation of consequences across time.

Table 3: Dimensions of transition justice in a near-critical political field.

Dimension	Diagnostic question	Near-critical relevance
Distributional incidence	Which actors receive benefits, burdens, protection, exposure, or loss?	Small interventions can generate large and uneven downstream effects.
Agency and voice	Which affected actors can interpret, contest, revise, or refuse the transition?	Rapid trajectory selection can narrow opportunities for participation before consequences stabilize.
Political legitimacy	Through what authority, procedure, or public justification is leverage exercised?	High leverage can tempt effective steering before sufficient justification is established.
Option-space effects	Whose practically reachable futures expand or contract?	Trajectory closure can distribute future possibilities asymmetrically.
Reversibility	Which commitments can later be revised, compensated, or exited?	Hysteresis and self-reinforcement can make correction substantially more costly after transition.
Temporal incidence	Which consequences are transferred to later actors or generations?	Persistent transitions can impose conditions on persons who did not participate in the initiating decision.

The definition is intentionally plural. It does not identify justice with one metric, nor does it presume that all dimensions can be aggregated without remainder. Its role is to prevent dynamical descriptions such as "stable," "efficient," "high-leverage," or "aligned with the prevailing propensity" from becoming substitutes for normative evaluation.

Proposition 12.2 (State–path separation)**.** *The normative assessment of a political transition requires separate attention to the resulting configuration and to the process that produces it. Evidence that one configuration performs better according to a selected outcome measure supplies insufficient information about the justice or legitimacy of the path through which that configuration is reached.*

This separation also clarifies the limits of attractor language. An attractor may be dynamically persistent while distributing security, recognition, resources, or political standing unevenly. A politically desirable property in one dimension can coexist with serious loss in another. The vocabulary of attractors therefore describes persistence and reachability; it does not carry an intrinsic ranking of political worth.

Table 3 organizes the dimensions used in the remainder of the section. The table is diagnostic rather than algorithmic. It provides a structure for asking what a transition redistributes when the redistributed objects include political possibilities as well as material outcomes.

12.2 Distributional incidence and normative amplification

Near criticality can create a gap between the physical, administrative, or organizational magnitude of an intervention and the scale of its normative consequences. Let $b_i(\gamma)$ denote the bundle of benefits and burdens borne by actor or group i along transition path γ. A near-critical field may permit

$$\|u_t\| \ll 1, \qquad \|\Delta b_i(\gamma)\| \gg \|u_t\| \quad \text{for some } i. \tag{57}$$

The expression is illustrative: it does not assume commensurability among all benefits and burdens. Its purpose is to mark a political asymmetry. Low intervention cost for the decision maker can coexist with high consequence for those situated downstream of the transition.

Definition 12.3 (Normative amplification)**.** *Normative amplification* describes a condition in which elevated susceptibility allows a comparatively limited intervention to contribute to large changes in the distribution of benefits, burdens, risks, standing, or future possibilities. The concept concerns the scale

of normative consequence relative to initiating leverage rather than the moral quality of the resulting change.

Normative amplification creates a special danger for control-oriented reasoning. If an intervention is attractive because it is inexpensive, timely, and dynamically effective, those same properties can obscure the scale of loss experienced by actors whose interests are weakly represented in the decision process. The problem is therefore broader than conventional cost–benefit comparison. Distributive justice concerns which benefits and burdens are politically produced and how competing distributive principles evaluate those allocations (Lamont and Favor 2017). Near criticality adds the question of whether a leverage-rich transition changes the future distribution of political capabilities themselves.

Working Claim 12.4 (Leverage–justice independence)**.** High dynamical leverage supplies information about efficacy. It does not by itself establish the justice, legitimacy, or ethical permissibility of using that leverage. As susceptibility rises, the burden of distinguishing these judgments becomes more important because small interventions can contribute to persistent distributive consequences.

This claim also applies when the intervention aims at an outcome widely regarded as beneficial. A transition can reduce one serious harm while concentrating new risks on a smaller group. It can improve aggregate welfare while reducing political standing or exit options for particular actors. It can protect a majority while making a minority bear the irreversibility of the transition. The framework therefore resists any inference from aggregate improvement to a complete judgment of transition justice.

12.3 Agency, participation, and transition standing

The temporal structure of near-criticality raises a question of standing: who is entitled to participate in decisions that can help close a collective trajectory? Legitimacy theory offers several competing answers to the broader problem of justified political power, including accounts centered on consent, public justification, participation, fair procedure, substantive values, or combinations of these (Peter 2023). The present framework does not select one theory. It isolates a condition under which the legitimacy question becomes especially acute.

A high-susceptibility field may compress decision time. Actors who identify a narrow window of leverage can argue that consultation, contestation, or procedural delay will allow the system to cross a threshold. This possibility is real in some settings, yet it can also become a recurrent justification for reducing participation precisely when the consequences of intervention are largest. Near-criticality therefore produces a tension between temporal responsiveness and political inclusion.

Let $V_i(\gamma)$ denote the effective capacity of actor i to register a claim, contest a description of the situation, influence the response, or seek revision during transition path γ. A justice analysis should attend to changes in V_i alongside changes in material outcomes:

$$\Delta V_i = V_i(\gamma_{\text{after}}) - V_i(\gamma_{\text{before}}). \tag{58}$$

Again, the notation is conceptual rather than a promise of a universal metric. Its purpose is to keep agency visible when institutional outcomes dominate the analysis.

Definition 12.5 (Transition standing)**.** *Transition standing* refers to the normative claim of affected actors to participate in, contest, receive justification for, or otherwise bear political relation to decisions that can substantially shape a shared trajectory. The content of such standing depends on the underlying theory of legitimacy and justice.

The definition allows degrees and forms of participation. Some near-critical contexts may require delegated authority, rapid protective action, or temporary constraint. Those conditions shift the justificatory burden rather than eliminating it. A response architecture can therefore be assessed partly by whether

urgency is paired with later review, explanation, compensation, restoration of voice, or reopening of decisions that were temporarily closed.

Research Obligation 12.6 (Standing audit)**.** When near-critical leverage is used to accelerate a political transition, analysis should identify which affected actors participated in the decision, which actors were represented indirectly, which actors had limited opportunity to contest the diagnosis, and which channels remain available for later review or revision.

12.4 Option-space distribution

Section 5 introduced trajectory closure as a contraction of practically reachable futures. The justice question concerns the distribution of that contraction. Let $\Omega_i(t)$ denote the set of future configurations that remain meaningfully reachable for actor or group i under current institutional, material, and political constraints. A transition can alter these sets unevenly:

$$\Delta\Omega_i = \Omega_i(t+1) - \Omega_i(t), \qquad \Delta\Omega_i \neq \Delta\Omega_j. \tag{59}$$

The set difference is schematic. Political options differ in quality, accessibility, and meaning, and those differences are rarely reducible to a simple count. The central point is distributional: an intervention can preserve collective optionality in the aggregate while sharply contracting the futures available to specific actors.

This issue extends the asymmetry test introduced in Section 10. Restraint can preserve institutional flexibility by transferring risk to actors who lack the resources to wait. Intervention can preserve one group's security by reducing another group's capacity to exit, dissent, organize, or recover. A trajectory that appears "open" at the system level can therefore be experienced as closure at the actor level.

Definition 12.7 (Asymmetric option contraction)**.** *Asymmetric option contraction* occurs when a political transition reduces the practically reachable future alternatives of some actors substantially more than those of others, especially when the affected actors have limited standing in the transition process or limited capacity to reverse the loss later.

The concept does not assume that every contraction of options is unjust. Political orders routinely constrain some courses of action in order to secure rights, public goods, or protection from harm. The concept instead identifies a normative object that becomes easy to miss when the transition is described only through system-level attractors. It asks whose possibility space is being used as the material through which collective stability or adaptation is purchased.

12.5 Temporal incidence and future actors

Persistent transitions distribute consequences across time as well as across contemporaries. Intergenerational justice addresses the asymmetric relation between present decision makers and future persons who can be affected by current institutional and resource choices without participating in them (Meyer 2021). Near-critical transitions sharpen this problem when current action changes the set of institutional, ecological, or political configurations that later actors inherit.

The relevant temporal object is broader than a discounted stream of benefits and costs. Some transitions can alter the background conditions under which future actors make choices. This includes inherited infrastructure, legal and administrative capacities, resource constraints, political coalitions, security structures, and the presence or absence of viable alternatives. Work on irreversible loss likewise identifies option preservation under uncertainty as one reason to avoid some irreversible changes (Spiekermann 2022).

Let $\Omega_g(t+k)$ denote the practical option space available to a later cohort or generation g. A present transition can produce

$$\Omega_g(t+k \mid \gamma_A) \neq \Omega_g(t+k \mid \gamma_B), \tag{60}$$

while the members of g remain absent from the decision that selects between γ_A and γ_B. This asymmetry makes temporal incidence a component of transition justice even when the identities of future persons cannot be specified in advance.

Working Claim 12.8 (Temporal extension of trajectory closure)**.** When a near-critical transition persistently changes the institutional or material conditions inherited by later actors, trajectory closure acquires an intergenerational dimension. The normative assessment of the transition should therefore include the distribution of future option spaces alongside the interests of present participants.

This claim remains compatible with disagreement about the precise rights of future persons, discounting, population ethics, and non-identity. Its more limited function is to prevent persistence from being treated as a purely technical virtue. Persistence can be beneficial, burdensome, or mixed; its normative evaluation depends partly on what later actors inherit and how much capacity for revision remains available to them.

12.6 Transition justice under uncertainty

Near-critical justice decisions occur under several uncertainties at once. Political actors can be uncertain about the system's susceptibility, the probability of competing transitions, the identities of those most affected, the persistence of the resulting regime, and the moral principles that should govern tradeoffs among claims. Moral decision-making under uncertainty becomes particularly difficult when familiar ethical thought experiments remove the empirical uncertainty that characterizes real political action (Tarsney et al. 2024). The trolley analogy discussed informally in the motivation for this paper therefore becomes less determinate in a near-critical field: the tracks can move, the number and identity of affected actors may be uncertain, and pulling the lever can change the structure of later choices.

The framework should consequently avoid a single scalar "justice score." A more cautious representation is a transition-justice profile

$$\mathbf{J}(\gamma) = \big(D(\gamma), V(\gamma), L(\gamma), O(\gamma), R(\gamma), T(\gamma)\big), \tag{61}$$

where D represents distributional incidence, V agency and voice, L legitimacy conditions, O option-space effects, R reversibility, and T temporal incidence. The vector preserves plurality among dimensions. Different theories of justice can assign different priorities, thresholds, or relations among these components.

Proposition 12.9 (Transition-justice plurality)**.** *Near-critical political analysis should preserve the plurality of normative dimensions that can be affected by trajectory selection. Dynamical stability, aggregate welfare, procedural legitimacy, equal standing, option preservation, and reversibility can support different judgments in the same case; the framework therefore treats their relation as an object of political reasoning rather than prespecifying a universal aggregation rule.*

The proposition has an important practical implication. A policymaker who can identify an apparently attractive attractor has completed only one part of the analysis. The next questions concern who defines attractiveness, which claims are omitted by the selected metric, whether the transition path preserves political standing, and whether those who bear irreversible loss possess meaningful recourse.

Research Obligation 12.10 (Transition-justice record)**.** Near-critical interventions with potentially persistent effects should be accompanied, where feasible, by an explicit record of the expected distribution

of benefits and burdens, affected actors' standing, option-space effects, reversibility assumptions, temporal incidence, major uncertainties, and the reasons used to justify trajectory-shaping action. Such a record does not settle the normative dispute; it makes the structure of the dispute available for later scrutiny and revision.

The argument of this section therefore changes the meaning of "successful" near-critical response. Success cannot be inferred solely from reaching a stable configuration, exploiting a favorable propensity, or avoiding one identified hazard. A transition can be dynamically effective while remaining normatively contested, and its deepest costs may appear through the options it closes rather than the immediate losses it produces. The next section develops one dimension of this problem in greater depth: the ethics of irreversibility and risk when acting, delaying, and refraining can each carry persistent consequences.

13 Ethics of Irreversibility and Risk

This section develops the ethical dimension of near-critical response when the consequences of action can persist beyond the interval in which the original reasons for action remain valid. Its structural role is to extend the justice analysis of Section 12 from the distribution of transition effects to the ethics of acting under uncertain and potentially asymmetric irreversibility. The discussion proceeds through five connected objects: persistence, risk imposition, reversibility, delay, and moral uncertainty. The aim is to clarify the ethical structure of near-critical response without reducing that structure to a single precautionary rule.

Near criticality, ethical evaluation becomes difficult partly because the relevant probabilities can be uncertain, the transition mechanism can be only partly understood, and the effects of a response can alter the conditions under which later correction would occur. Philosophical work on risk emphasizes that risk assessment concerns both knowledge and lack of knowledge, and that real social decisions commonly involve probability estimates that remain uncertain, alongside cases that can be represented as more fully specified lotteries (Hansson 2022). Moral decision-making under uncertainty adds a further layer: ethical judgment can remain unsettled about the long-run consequences, the admissibility of risk imposition, and the relative force of competing moral principles (Tarsney et al. 2024). Near-critical politics combines these uncertainties with a dynamical feature: some mistakes can become harder to correct after the system has moved into a new regime.

Table 4 places action, delay, buffering, and restraint in the same ethical field. The table is intentionally non-ranking. Each mode can preserve one value while placing another at risk. This symmetry is important because political discussions often assign ethical responsibility to visible interventions while treating delay or restraint as comparatively neutral. Near a critical region, the temporal evolution of the system makes that distinction difficult to sustain in general form.

13.1 Irreversibility as an ethical multiplier

Irreversibility matters ethically because later correction can lose practical meaning even when retrospective learning remains possible. Spiekermann's account of irreversible loss identifies several reasons for preserving options under uncertainty, including prudential option value and obligations to future persons (Spiekermann 2022). The present framework extends that concern to political transitions in which institutions, expectations, coalitions, or material infrastructures become self-reinforcing after a near-critical shift.

Let a denote a response selected at time t, and let $\Gamma(a)$ denote the set of trajectories that remain practically reachable after that response. A simple representation of reversibility can be written as

$$\mathcal{R}(a,t) = \Pr(x_t \in \mathcal{V} \mid a, \text{ later corrective action}), \tag{62}$$

where $\mathcal{V}$ is a reference region that retains some previously available institutional or political capacities. Equation (62) is schematic. The point is to distinguish an intervention's immediate effect from the practical possibility of revising its consequences later.

The ethical significance of a response can therefore increase with both the magnitude of possible harm and the difficulty of revision. One conceptual representation is

$$\mathcal{B}_{\text{irr}}(a) = \mathcal{H}(a) \cdot \mathcal{P}(a) \cdot \big(1 - \mathcal{R}(a)\big), \tag{63}$$

where $\mathcal{H}$ denotes the severity of affected interests, $\mathcal{P}$ the plausibility of the relevant transition, and $\mathcal{R}$ practical reversibility. The expression is illustrative rather than a proposal for universal ethical measurement. Its structural function is to keep persistence visible when a response appears attractive because its initial cost is small.

Table 4: Ethical dimensions of response under possible near-critical transition.

Response mode	**Potential ethical value**	**Characteristic risk**	**Reversibility concern**	**Distributional question**
Active intervention	Protection, interruption of harmful cascade, preservation of vulnerable interests	Misdiagnosis, amplification, trajectory selection beyond mandate	Policy may trigger or deepen a regime shift before evidence is revised	Who bears losses created by the intervention and who authorized exposure?
Delay and observation	Additional information, reduced premature closure, opportunity for distributed judgment	Waiting can permit harmful feedback or option loss to accumulate	Time itself can move the system across a threshold	Who can afford to wait and whose options deteriorate during observation?
Buffering and reversible response	Harm reduction with retained room for correction	Temporary measures can still redistribute burdens or become politically entrenched	Reversal may be easier in design than in practice	Which groups carry the costs of maintaining flexibility?
Restraint or non-forcing	Avoidance of unnecessary steering, preservation of emergent adaptation	Existing power and endogenous dynamics continue operating	Non-intervention can allow irreversible closure through uncontrolled processes	Whose status quo is preserved and whose exposure remains unaddressed?

Definition 13.1 (Irreversibility burden)**.** The *irreversibility burden* of a near-critical response is the additional ethical weight generated when the response can contribute to persistent loss, trajectory closure, or institutional transformation that later actors may find substantially harder to revise than the original intervention was to initiate.

Working Claim 13.2 (Low-force, high-burden possibility)**.** A near-critical intervention can require little material or institutional force while carrying a substantial irreversibility burden. Dynamical efficiency and ethical burden can therefore vary independently.

This possibility is central to the paper's rejection of a simple control orientation. The availability of a low-cost perturbation leaves the ethical significance of the transition it helps produce. In some cases, high susceptibility can increase the responsibility associated with small actions because the system converts limited input into persistent collective change.

13.2 Risk imposition and political standing

The ethics of risk concerns more than the probability-weighted value of final outcomes. Hansson emphasizes distinctions among voluntary risk-taking, accepted risk exposure, and risks imposed on persons without their acceptance (Hansson 2022). These distinctions are especially relevant near criticality because an actor may intentionally exploit a leverage-rich configuration while the affected population remains uncertain about both the transition and the consequences of the intervention.

Let $\rho_i(a)$ denote the risk exposure imposed on actor or group i by response a, broadly understood to include material harm, loss of standing, loss of exit, and contraction of future options. An aggregate quantity such as $\sum_i \rho_i(a)$ captures only part of the ethical problem. Assessment also depends on how exposure is distributed, whether the exposed actors possess transition standing, whether they benefit from reciprocal arrangements, and whether compensation or later revision remains available.

Definition 13.3 (Transition risk imposition)**.** *Transition risk imposition* occurs when an intentional political response changes another actor's exposure to a potentially persistent regime change, including exposure created through altered feedback, institutional lock-in, or contraction of practical future options.

The definition includes indirect exposure. A response can leave an actor's formal rights unchanged while changing the surrounding field so that those rights become harder to exercise. It can also protect one population from an identified hazard by shifting uncertainty toward another population with less voice. Near-critical ethics therefore requires attention to the relational structure through which risk moves, alongside the actor that initiates the response.

Research Obligation 13.4 (Risk-imposition record)**.** When a near-critical response intentionally alters the probability or persistence of a political transition, the justification should identify, where feasible, the actors whose risk exposure is increased, the actors whose risk is reduced, the channels through which exposure is transferred, and the forms of standing, review, or compensation available to those carrying the additional risk.

This obligation remains relevant under urgency. Time pressure can narrow the feasible forms of consultation, while the distributional fact remains that some actors can be made to carry the consequences of a decision taken for collective protection. Recording that structure supports later accountability and helps distinguish emergency compression of procedure from permanent political exclusion.

13.3 Reversibility and response design

The value of reversibility becomes especially salient when the system is susceptible and knowledge remains provisional. A reversible response can create a period in which consequences are observed before stronger commitment occurs. This idea resembles the option-preserving rationale discussed in work on irreversible loss (Spiekermann 2022), while the present paper applies it to the design of political response sequences.

Let $a_1, a_2, \ldots, a_m$ denote responses of increasing closure pressure. A response architecture can preserve a sequence

$$a_1 \prec a_2 \prec \cdots \prec a_m, \tag{64}$$

where each stage is followed by observation, revision of the susceptibility assessment, and reconsideration of normative warrant. The ordering $\prec$ indicates increasing commitment or difficulty of reversal; moral ranking remains a separate operation.

Definition 13.5 (Reversibility premium)**.** A *reversibility premium* is the additional ethical value that a response can acquire, under substantial uncertainty, by preserving credible pathways for correction, withdrawal, compensation, or reopening of political choice.

The premium is conditional. Reversibility can lose value when delay itself causes severe harm, when a temporary response transfers large burdens to actors with little protection, or when nominally reversible measures generate expectations and dependencies that make withdrawal politically costly. The design question is therefore broader than legal sunset clauses. It concerns the actual social and institutional capacity to revise course.

Working Claim 13.6 (Reversibility–urgency trade-off)**.** Near criticality can create a conflict between preserving reversibility and responding rapidly enough to prevent severe trajectory closure. Ethical analysis should treat the trade-off as configuration-specific, with caution and speed receiving weight from the structure of the case.

This claim places the paper at some distance from a simple precautionary formula. Precautionary reasoning can support protective measures when evidence of serious danger remains incomplete, while the philosophy of risk also warns that practical decisions involve choices about false positives, false negatives, and the treatment of uncertain probabilities (Hansson 2022). Near criticality adds the possibility that both kinds of error can become persistent: premature intervention can induce a harmful transition, while premature restraint can allow an endogenous cascade to cross a threshold.

13.4 Delay, restraint, and temporal responsibility

Delay is often described as the absence of a decision, yet in a dynamical system it is itself temporally situated. If the underlying configuration continues to evolve, waiting changes the state at which the next decision is made. The ethical structure can be represented as

$$x_{t+\Delta t} = \Phi_{\Delta t}(x_t, \eta_{t:t+\Delta t}), \tag{65}$$

where $\Phi_{\Delta t}$ denotes the uncontrolled or partially governed evolution of the system during the waiting interval and η the perturbations entering that interval. A later response therefore confronts a different configuration.

This observation gives delay an ethical profile of its own. Waiting can create valuable information, reduce premature closure, and permit wider participation. It can also consume the time available for low-force response, permit harmful feedback to deepen, or transfer risk onto actors whose losses accumulate during the observation period. Section 10 introduced an asymmetry test for restraint. The present section extends that test toward temporal responsibility.

Definition 13.7 (Temporal risk transfer)**.** *Temporal risk transfer* occurs when the decision to delay or refrain from intervention preserves flexibility for some actors or institutions by shifting increasing exposure, option loss, or recovery cost onto others during the waiting period.

The concept prevents a simple ethical contrast between coercive action and innocent inaction. It also prevents the opposite mistake of treating every delay as culpable. A temporally responsible response requires attention to the rate at which different harms, uncertainties, and political options change during the waiting interval.

Let $c_i(t)$ denote the expected rate of option contraction or accumulating harm for actor i. A waiting interval Δt has an ethical significance that partly depends on the distribution

$$\left\{ \int_t^{t+\Delta t} c_i(s)\, ds \right\}_{i=1}^{n}. \tag{66}$$

This notation is again conceptual. It makes visible that one institution's "time to learn" can be another actor's period of deteriorating safety or political agency.

Research Obligation 13.8 (Delay audit)**.** Where restraint or extended observation is chosen near a suspected critical region, analysis should record which uncertainties are expected to decrease, which harms can accumulate during the interval, whose option spaces are likely to contract, and which indicators would justify shortening or ending the waiting period.

The audit makes restraint an accountable response with its own evidential and ethical record. It is consistent with the paper's treatment of *wuwei* as responsive participation: non-forcing can remain active through observation, buffering, protection of vulnerable actors, and preparation for revision even when trajectory steering is withheld.

13.5 Deep uncertainty and ethical decision structure

Near-critical politics can involve uncertainty at several levels. Actors may be uncertain about the current state, the location of a transition region, the noise process, the effects of intervention, the identities of future beneficiaries and losers, and the moral principles governing trade-offs. Tarsney, Thomas, and MacAskill distinguish several dimensions of moral decision under uncertainty and emphasize the difficulty of extending idealized ethical judgments into conditions of incomplete information (Tarsney et al. 2024). The philosophy of risk likewise distinguishes decisions under idealized known probabilities from real decisions in which probability estimates themselves remain uncertain (Hansson 2022).

A near-critical response problem can therefore be represented by a layered uncertainty profile

$$\mathbf{U} = \big(U_x, U_\theta, U_\eta, U_a, U_T, U_N\big), \tag{67}$$

where U_x denotes uncertainty about the present configuration, U_θ uncertainty about structural parameters or couplings, U_η uncertainty about noise, U_a uncertainty about intervention effects, U_T uncertainty about long-run trajectories, and U_N normative uncertainty. The vector preserves heterogeneity among uncertainties and keeps them visible beyond a single confidence number.

Proposition 13.9 (Uncertainty-sensitive ethical burden)**.** *The ethical burden of near-critical response depends partly on the structure of uncertainty, the persistence of possible consequences, and the distribution of risk exposure. A low-confidence estimate about a reversible outcome and a low-confidence estimate about an irreversible loss therefore present different ethical decision structures even when their numerical uncertainty appears similar.*

This proposition supports a differentiated form of caution. Some cases may justify rapid protective action because the cost of a missed transition is severe and later correction is unlikely. Other cases may justify delay because the principal hazard lies in premature trajectory closure and the cost of waiting remains bounded. A third class may support buffering measures that protect vulnerable actors while preserving several later choices. The paper therefore treats precaution as a family of response architectures whose form depends on the configuration.

13.6 Ethical hysteresis and the persistence of responsibility

Dynamical hysteresis concerns differences between the conditions of transition and the conditions of return. Political responsibility can display an analogous asymmetry without being reducible to the physical concept. A decision can be easier to initiate than to undo, and the duties generated by the decision can persist after the original decision makers lose office, information changes, or the initial emergency ends.

This paper uses the phrase *ethical hysteresis* cautiously for that institutional asymmetry. The phrase marks a limited analytical pattern. It identifies cases in which a response creates enduring obligations of repair, compensation, review, or restoration that survive the temporary conditions that originally justified the action.

Definition 13.10 (Ethical hysteresis)**.** *Ethical hysteresis* refers to a persistence of normative responsibility after the initiating conditions of a near-critical response have changed, especially where the response has created durable burdens, dependencies, lost options, or institutional asymmetries that later actors inherit.

The concept is useful because emergency reasoning often concentrates on the moment of authorization. Near-critical ethics also requires attention to the post-transition interval. A temporary intervention can produce durable losers. A protective restriction can create new institutional dependencies. A delayed response can leave harms that remain after the danger has passed. The ethics of response therefore includes the obligations that emerge from the history of the transition itself.

Working Claim 13.11 (Persistence of response responsibility)**.** Where a near-critical response contributes to persistent burdens or option loss, the justificatory responsibility associated with that response can extend beyond the period of immediate intervention. Review, restoration, compensation, or reopening of political choice can remain ethically relevant after the critical interval has ended.

This temporal extension also connects ethics with trajectory records. If later actors are to evaluate whether emergency leverage was used responsibly, the political system needs records of the evidence, uncertainty, expected risks, standing of affected groups, reversibility assumptions, and reasons for action or restraint. The transition-justice record proposed in Section 12 therefore also functions as an ethical memory mechanism.

13.7 A provisional ethic of near-critical response

The preceding analysis yields a structured set of ethical considerations that can organize judgment while leaving the choice of any universal decision rule open to substantive moral argument. A near-critical response can be examined through at least six questions: the severity and distribution of the risk, the quality and structure of uncertainty, the practical reversibility of the response, the rate at which delay changes exposure, the standing of affected actors, and the persistence of obligations created by the response.

These considerations can be summarized schematically as

$$\mathcal{E}(a) = \mathcal{F}(\mathbf{J}(a), \mathbf{U}(a), \mathcal{R}(a), \mathbf{C}(a), \mathbf{V}(a), \mathbf{M}(a)), \tag{68}$$

where **J** represents the transition-justice profile introduced in Section 12, **U** uncertainty, $\mathcal{R}$ reversibility, **C** temporal cost or option contraction, **V** standing and voice, and **M** the obligations of memory, repair, or review. The function $\mathcal{F}$ is intentionally unspecified. Different ethical and political theories can order these dimensions differently.

Proposition 13.12 (Ethics before trajectory commitment)**.** *When a political system appears near a sensitive transition, information about leverage should be accompanied by an explicit ethical assessment of risk imposition, reversibility, delay, standing, and persistence before intentional trajectory commitment is treated as warranted. The proposition supplies a procedural orientation while leaving substantive rankings of future attractors to explicit normative argument.*

This orientation preserves the distinction developed throughout the paper. Criticality analysis can identify sensitivity. Response analysis can identify possible modes of participation. Ethics asks how the burdens of action, restraint, uncertainty, and persistence are borne. These operations remain distinct, and each contributes information or judgment that the others leave open.

The next section turns to a different form of judgment. Near criticality, actors can sometimes perceive an emerging pattern before its consequences and normative status are fully manifest. Such perception can organize attention and inquiry, yet it can also aestheticize coherence, momentum, or order in ways that conceal exclusion and harm. The next section therefore examines aesthetic judgment and emerging form as a distinct, non-substitutive component of near-critical political reasoning.

14 Aesthetic Judgment and Emerging Form

This section introduces aesthetic judgment as a distinct mode of attention to near-critical political fields. The preceding sections separated dynamical propensity, efficacy, justice, legitimacy, risk, and reversibility. A further problem remains. Near a transition, possible forms of organization can become partially perceptible before they are fully manifest, causally understood, or normatively assessed. Political actors may sense coherence, tension, rhythm, fragility, dissonance, or an emerging pattern while the available concepts and models remain incomplete. The structural role of this section is to describe that pre-manifestational mode of judgment without granting it epistemic or normative authority that it does not possess.

The discussion proceeds in five stages. First, it distinguishes reflective attention to emerging form from determinate classification. Second, it defines *pre-manifestational aesthetic judgment* as a provisional recognition of form worthy of further attention. Third, it relates this judgment to noise, resistance, rhythm, and temporal development. Fourth, it examines plurality and communicability, since political form appears differently from different positions in the field. Fifth, it develops the danger of *aesthetic capture*: the conversion of perceived elegance, coherence, or momentum into an unearned claim of truth, justice, or political warrant.

The comparative philosophical resources used here have limited functions. Kant's account of reflective judgment is relevant because it distinguishes the activity of seeking a universal for a given particular from the application of an already supplied concept, and his aesthetics separates judgments of beauty from determinate judgments of the good (Ginsborg 2022). Dewey's aesthetics is relevant because it treats form and aesthetic experience as temporally developing relations among doing, undergoing, resistance, continuity, accumulation, and environment (Leddy and Puolakka 2025). Arendt's political appropriation of Kantian judgment is relevant to plurality, communicability, and the effort to consider a common world from perspectives beyond one's own (Tömmel and Passerin d'Entreves 2025). These traditions are used as comparative supports. The present concept of pre-manifestational judgment is a working construction for near-critical political analysis rather than an interpretive claim that any one of these authors already supplied it.

14.1 Incomplete manifestation and reflective attention

Near-critical fields can contain configurations whose future significance is not yet settled. A coalition may be forming without having stabilized. A local practice may be acquiring wider resonance without constituting a new regime. A pattern of refusal, adaptation, institutional learning, or public coordination may appear across otherwise disconnected events. The analyst can possess insufficient grounds for a determinate claim such as "this is the next attractor" while still having reasons to attend to the emerging relation.

Kant's distinction between determining and reflecting judgment offers a useful analogy. Determining judgment subsumes a particular under a concept already available, whereas reflecting judgment searches for an appropriate universal for a given particular (Ginsborg 2022). The present paper does not import Kant's transcendental framework into political dynamics. It adopts only the structural contrast between classification under a settled rule and judgment exercised while the relevant rule or concept remains incomplete.

Let z_t denote a partially manifested configuration observed at time t, and let $\mathcal{C}_t$ denote the set of concepts and models currently available to an observer. A determinate classification is possible when some $c \in \mathcal{C}_t$ adequately organizes the relevant features of z_t. Near criticality, the observer can instead confront

$$z_t \not\models c \quad \text{adequately for any available } c \in \mathcal{C}_t, \tag{69}$$

while still perceiving relational organization in the configuration. Equation (69) represents incomplete

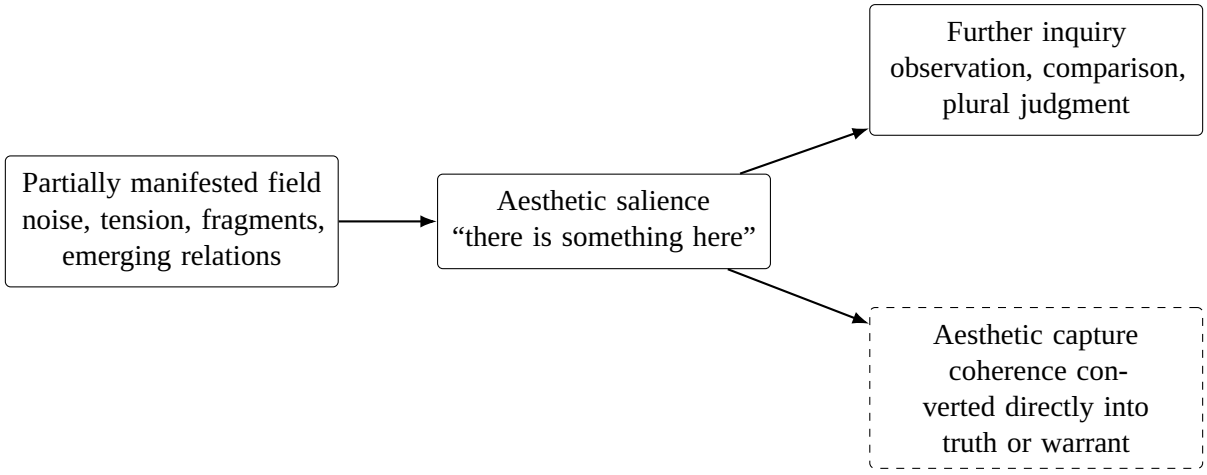

Figure 9: Aesthetic salience as an intermediate stage between incomplete manifestation and subsequent inquiry. The lower branch marks the risk of converting perceived form directly into epistemic or normative authority.

conceptual fit. It leaves open whether the apparent organization will persist, disappear, or be re-described under a later model.

Definition 14.1 (Pre-manifestational aesthetic judgment)**.** *Pre-manifestational aesthetic judgment* is a provisional judgment that an incompletely manifested configuration exhibits a form, tension, coherence, rhythm, contrast, or generative potential worthy of sustained attention before its causal structure, long-run trajectory, or normative status can be adequately determined.

The term *aesthetic* refers here to sensitivity to form and relation. It does not restrict the judgment to beauty, pleasure, or fine art. It also does not imply that political judgment should be governed by taste. The concept is introduced because near-critical analysis confronts situations in which the observer must decide what deserves further attention before a complete propositional account is available.

Figure 9 places aesthetic judgment in an intermediate position. Its primary function is attentional: it can identify a configuration that merits further inquiry. The upper branch preserves this provisional status through observation, comparison, and plural judgment. The lower branch represents a recurrent political danger. A pattern that appears coherent, elegant, harmonious, inevitable, or historically powerful can be converted into a claim that the pattern is true, desirable, or entitled to political realization. The remainder of the section develops reasons to keep those operations separate.

Working Claim 14.2 (Attention without authorization)**.** Pre-manifestational aesthetic judgment can supply a reason to attend more closely to an emerging configuration. It supplies no independent authorization to amplify, suppress, or institutionalize that configuration.

This claim gives aesthetic judgment a limited yet meaningful place in the response architecture. Near criticality creates a surplus of possible signals. Some fluctuations are measurement error, some are transient noise, some reveal susceptibility, and some participate in emerging organization. Formal models and historical evidence remain central to distinguishing among these cases. Aesthetic judgment enters earlier, at the point where an observer notices that one relation among many deserves the expenditure of scarce epistemic attention.

14.2 Temporal form, resistance, and noise

The relation between aesthetic judgment and near-criticality becomes clearer when form is treated temporally. Dewey’s account of aesthetic experience places strong emphasis on development, accumulation,

rhythm, resistance, and the interaction between an organism and its environment. On that account, resistance can participate in the development of an experience rather than functioning only as an obstacle, and form concerns the active relations among parts across time (Leddy and Puolakka 2025). The present framework adopts this temporal emphasis without treating political processes as artworks.

Let a candidate political form be represented provisionally by a sequence

$$\mathcal{F}_{t:t+k} = \{z_t, z_{t+1}, \ldots, z_{t+k}\}. \tag{70}$$

The form is then judged through relations among phases rather than through one isolated state. Features that can become salient include continuity, recurrent contrast, accumulation, interruption, mutual adjustment, and the way later phases transform the significance of earlier ones. These features are neither necessary nor sufficient conditions for political importance. They specify the kind of temporal organization to which aesthetic attention can be directed.

This temporal view also changes the treatment of noise. Section 3 argued that noise can obscure structure, reveal susceptibility, induce transition, or participate in emerging organization. Aesthetic judgment adds a related epistemic possibility: irregularity can contain incipient form before the form is expressible within the prevailing model. The observer may register an unusual rhythm, recurring juxtaposition, or patterned tension before a causal account is available.

The same capacity creates a symmetric danger. Human observers are highly capable of perceiving patterns in stochastic variation. Aesthetic salience can therefore generate false positives. The experience that "these events belong together" remains compatible with later evidence that they do not. Near criticality intensifies the stakes because a prematurely interpreted pattern can guide interventions that alter the field itself.

Proposition 14.3 (Noise–form ambiguity)**.** *In a near-critical field, perceived form in noisy observations is compatible with at least three possibilities: an emerging relational structure, a transient configuration, or an observer-dependent pattern imposed on weakly related events. Aesthetic salience therefore supports inquiry under ambiguity rather than a direct inference to underlying order.*

The proposition strengthens the paper's earlier injunction to epistem the noise. Epistemic work includes learning how a perceived form behaves under new observations, whether it survives changes of scale, whether multiple observers recognize related organization, and whether the form produces consequences consistent with the proposed relational account. Aesthetic recognition can begin this process. It cannot complete it.

14.3 Form salience and observer position

Aesthetic judgment is situated. What appears as coherent form from one position can appear as rupture, exclusion, or disorder from another. This is especially important in political systems because actors occupy different locations in networks of risk, authority, recognition, mobility, and dependence. An emerging configuration that appears elegantly coordinated to a central institution can appear as narrowing option space to actors at its margins.

Let o_i denote an observer situated at position i, with historical experience H_i and access to observations O_i. Aesthetic salience can be represented schematically as

$$\Sigma_i(z_t) = \Sigma(z_t; H_i, O_i), \tag{71}$$

where Σ_i is a salience profile rather than a scalar score. Different observers can register different forms because their histories and positions make different relations available to perception.

This situatedness gives political relevance to communicability. Kant's theory of taste links aesthetic judgment to a claim that seeks agreement without reducing the judgment to a determinate concept, while

Arendt's political use of Kantian judgment emphasizes the effort to consider phenomena from perspectives beyond one's private standpoint and the non-coercive character of appealing to others' judgment (Ginsborg 2022; Tömmel and Passerin d'Entreves 2025). The present paper adopts a narrower implication: when aesthetic salience influences political attention, its public articulation should remain open to perspective change and contestation.

Definition 14.4 (Plural form recognition)**.** *Plural form recognition* is the practice of comparing how an emerging configuration is perceived from different political and social positions before its apparent coherence is treated as a stable description of the field.

Plural form recognition is especially valuable when susceptibility is high. One observer may detect coordination that others experience as coercive synchronization. Another may identify noise where participants experience a newly meaningful relation. A third may recognize costs or exclusions hidden by the dominant representation. The aim is not to average these perceptions into one aesthetic consensus. It is to expose the dependence of form recognition on position and to prevent a privileged standpoint from becoming invisible.

Research Obligation 14.5 (Perspective record)**.** When perceived form is used to allocate political attention near criticality, analysis should record, where feasible, whose perspective produced the judgment, which actors or relations are visible from that position, which remain weakly observed, and how alternative situated judgments change the description of the emerging configuration.

This obligation links aesthetics to the justice analysis of Section 12. It does not convert aesthetic plurality into a theory of justice. Its purpose is antecedent: to reduce the probability that an emerging order is aesthetically represented through the experience of its beneficiaries while the burdens experienced elsewhere remain classified as disorder or noise.

14.4 Aesthetic capture and normative separation

Political history supplies many reasons to treat aesthetic coherence with caution. Order, unity, symmetry, scale, speed, technological sophistication, and collective synchronization can generate powerful impressions without resolving questions of justice or legitimacy. A configuration can appear coherent because opposition has been excluded. It can appear stable because exit has become costly. It can appear elegant because heterogeneous experiences have been compressed into a single representation.

The political relation between aesthetics and visibility also has a broader philosophical literature. Rancière, for example, connects aesthetics and politics through the distribution of what becomes visible, audible, thinkable, and possible (Rancière 2013). The present framework does not adopt his theory as a whole. It draws one caution from this connection: what is available to aesthetic judgment is already shaped by institutions, categories, media, and distributions of attention. Aesthetic perception therefore cannot be presumed politically innocent.

Definition 14.6 (Aesthetic capture)**.** *Aesthetic capture* occurs when the perceived coherence, elegance, intensity, harmony, momentum, or historical drama of an emerging configuration is converted into epistemic certainty, normative superiority, or political warrant without completing the relevant evidential, ethical, and legitimacy assessments.

Aesthetic capture can operate in several directions. A disruptive movement can be romanticized because turbulence appears generative. A highly ordered institution can be privileged because coherence appears efficient. A rapid technological transition can be celebrated because acceleration appears historically inevitable. A decentralized configuration can be valued because spontaneity appears ethically attractive. Each conversion replaces a distinct argument with an aesthetic impression.

Proposition 14.7 (Aesthetic–normative separation)**.** *The perception of emerging form and the evaluation of political worth are analytically distinct. Aesthetic salience can coexist with injustice, illegitimacy, high irreversible risk, or severe contraction of another actor's option space.*

The proposition parallels the leverage–justice independence developed in Section 12. Configuration-dependent propensity does not establish political worth; aesthetic coherence does not establish it either. Near-critical response therefore requires at least four separate operations: recognition of form, assessment of dynamical plausibility, evaluation of justice and ethics, and consideration of legitimate authority. The operations can inform one another while retaining different criteria.

Aesthetic judgment can also conflict with the ethics of reversibility. An emerging order may appear compelling precisely because it possesses strong momentum and internal coherence. Those properties can simultaneously indicate that later revision will become difficult. The aesthetic attraction of closure can therefore increase at the same moment that ethical reasons for preserving revisability become stronger.

Remark 14.8 (Beauty of closure)**.** A completed form often appears more coherent than an open field of unresolved possibilities. Near criticality, this experiential preference can create a bias toward premature trajectory closure. The paper treats openness, ambiguity, and unfinishedness as politically relevant conditions whose lower aesthetic coherence does not by itself count against them.

This remark matters for the relation between aesthetic judgment and *wuwei*. Non-coercive response can require tolerating an unfinished field whose eventual organization remains uncertain. The relevant aesthetic capacity is therefore not only appreciation of harmony. It can include sensitivity to productive tension, incompleteness, asymmetry, and forms still becoming visible.

14.5 Aesthetic judgment within near-critical response

The place of aesthetic judgment in near-critical response can now be stated more precisely. It is neither a substitute for evidence nor an independent source of political authorization. It is a mode of selective attention to emerging form under incomplete manifestation. Its practical value is greatest before strong claims are made, when scarce attention must be allocated among many weak, ambiguous, and potentially consequential signals.

A response architecture can therefore distinguish four stages:

(i) **aesthetic registration**: an observer notices a configuration, tension, rhythm, or relation that appears worthy of attention;

(ii) **epistemic development**: the candidate form is tested against additional evidence, alternative models, scales, and observations;

(iii) **plural and normative assessment**: affected perspectives, justice, legitimacy, risk, and reversibility are examined;

(iv) **response selection**: observation, buffering, redirection, restraint, or stronger intervention is considered within the broader response repertoire.

The ordering is analytical rather than chronologically rigid. Political actors can cycle among these stages as new information appears. Its purpose is to prevent the salience of an emerging form from collapsing the later stages into an immediate decision.

Working Claim 14.9 (Aesthetic contribution to near-critical inquiry)**.** Aesthetic judgment has a defensible role near criticality when it enlarges the set of potentially meaningful configurations available to inquiry while remaining revisable under evidence, plural perspectives, and normative assessment.

This role also clarifies why aesthetics belongs in a paper on near-critical response. Criticality creates conditions in which the fully manifested future cannot be observed directly, yet political actors cannot attend equally to all possible configurations. Some faculty of salience is unavoidable. Formal models, institutions, expertise, historical analogy, and embodied perception all participate in that allocation. Aesthetic judgment names one component of this process: sensitivity to relational form before that form is fully conceptualized.

The section therefore adds neither an aesthetic criterion of political order nor an instruction to follow whatever pattern appears compelling. Its contribution is a disciplined place for pre-manifestational judgment within a larger architecture of response. Form can be noticed before it is explained. Explanation can develop before intervention is warranted. A configuration can be aesthetically compelling while remaining unjust. An unfinished field can be politically valuable because its openness preserves possibilities that a more coherent order would close.

The next section extends this final point by examining option space, agency, and revisability directly. If near-critical politics concerns several still-reachable futures, then the distribution and preservation of those possibilities becomes an object of political analysis in its own right.

15 Option Space, Agency, and Revisability

This section develops the political significance of option space after the preceding analyses of trajectory closure, transition justice, irreversibility, and aesthetic judgment. Near a sensitive transition, a response can alter more than a presently observed distribution of benefits and burdens. It can also change which futures remain practically reachable, who retains meaningful capacity to choose among them, and whether collective decisions can later be reopened when information, preferences, or conditions change. The structural role of this section is therefore to treat option space, agency, and revisability as related but distinct objects of political analysis.

The discussion proceeds in six stages. First, it distinguishes abstract state space from politically meaningful option space. Second, it relates practical reachability to substantive opportunity and conversion conditions. Third, it examines agency and the security of options against unilateral closure by others. Fourth, it defines revisability as an institutional and relational capacity rather than a perpetual refusal of commitment. Fifth, it analyzes asymmetric option contraction across actors and generations. Sixth, it states a qualified option-preservation orientation whose force depends on justice, ongoing harm, and the costs of delay.

The capability approach provides one comparative vocabulary for this inquiry because it distinguishes resources or formal permissions from substantive opportunities to do and be things that people can actually realize (Robeyns and Byskov 2025). Republican theories of freedom provide a second comparative vocabulary by emphasizing the structural insecurity that arises when another actor retains uncontrolled capacity to interfere, even when interference is not presently exercised (Lovett 2026). Adaptive policy-pathway research supplies a more institutional comparison: it explicitly represents sequences of actions, lock-in, path dependence, and future adaptation under deep uncertainty (Haasnoot et al. 2013). The present framework does not derive a theory of justice from any of these literatures. It uses them to clarify why the existence, accessibility, security, and revisability of future possibilities should be kept analytically separate.

Table 5 emphasizes that the number of nominal alternatives is a poor proxy for political freedom. A field with ten formal choices can provide less meaningful agency than a field with three robust and substantively accessible pathways. Near-critical analysis therefore requires a richer object than cardinality alone.

15.1 Political option space and practical reachability

Let $\mathcal{X}$ denote a conceptual state space and let $\Omega_i(t) \subseteq \mathcal{X}$ denote the set of future configurations that remain practically reachable for actor or group i from time t, given the current institutional configuration, resources, dependencies, and expected transition costs. This construction is deliberately broader than a menu of explicit policy choices. A future can be physically conceivable while lying outside the practical option space of a particular actor.

One provisional representation is

$$\Omega_i(t;\tau) = \left\{ x' \in \mathcal{X} : C_i(x_t \rightsquigarrow x' \mid G_t, H_t) \leq \tau_i \right\}, \tag{72}$$

where C_i is a context-dependent transition cost, G_t the relational and institutional configuration, H_t relevant history, and τ_i a practical threshold beyond which a path ceases to be meaningfully available to actor i. Equation (72) is a conceptual device rather than an empirical metric. Its purpose is to make explicit that reachability is actor-relative, history-dependent, and conditioned by more than formal permission.

Definition 15.1 (Political option space)**.** The *political option space* of an actor or group is the set of future configurations and courses of action that remain substantively reachable under the prevailing distribution of resources, institutional permissions, relations, transition costs, and dependencies.

Table 5: Analytical dimensions of politically meaningful option space.

Dimension	Analytical concern	Near-critical relevance
Reachability	Whether a future configuration can be approached at feasible social, political, and material cost	A small transition can move previously accessible futures beyond practical reach
Substantive availability	Whether actors possess the resources, standing, and conversion conditions required to use an apparent option	Formal options can remain while practical agency contracts
Security	Whether an option remains available independently of another actor's uncontrolled discretion	Temporary openness can coexist with structural vulnerability to unilateral closure
Voice and participation	Whether affected actors can shape selection, sequencing, and revision of collective pathways	Trajectory choice can redistribute agency before final outcomes are known
Revisability	Whether an adopted pathway can later be altered, suspended, compensated, or reopened at tolerable cost	Hysteresis and lock-in can make later correction much harder than initial commitment
Temporal durability	How long options remain open and how present choices redistribute options across time	Present flexibility can be purchased through future constraint, or the reverse

This definition is compatible with the capability approach's emphasis on real or substantive opportunity while remaining narrower in purpose. Capability accounts examine what people are actually able to do and be, and they stress that resources alone do not determine real opportunity because conversion conditions differ across persons and environments (Robeyns and Byskov 2025). The present concept focuses on trajectory-sensitive political reachability. It asks how a near-critical transition changes the set of futures that actors can still meaningfully enter.

The distinction also clarifies the difference between present welfare and future openness. An intervention can improve a current outcome while making several later configurations prohibitively costly. Another intervention can preserve many formal paths while leaving affected actors without the resources or institutional standing required to use them. Option-space analysis therefore supplements outcome analysis rather than replacing it.

Proposition 15.2 (Outcome–option distinction)**.** *Evaluation of a near-critical transition can require separate attention to the resulting state and to the practical option spaces that remain available after the transition. Improvement in a presently observed outcome does not by itself establish preservation of future agency.*

This proposition extends the distinction between transition justice and state justice developed in Section 12. It also adds a temporal object to the ethics of irreversibility in Section 13: the ethical significance of irreversibility partly concerns the futures that become inaccessible after a commitment has been made.

15.2 Substantive availability and conversion conditions

An option can appear in a legal, technical, or institutional description while remaining practically unavailable. The capability literature offers a useful analytical reminder here: formal rights and resources are mediated by personal, social, and environmental conversion conditions before they become substantive freedoms (Robeyns and Byskov 2025). Near criticality, these conversion conditions can change quickly because institutions, prices, coalitions, infrastructure, and expectations may reorganize together.

For the present framework, let $a \in \Omega_i$ be an apparent option and let $\kappa_i(a, t)$ denote a vector of conversion conditions needed to make that option practically usable. A transition may leave the legal

description of a unchanged while altering κ_i so strongly that the option becomes nominal. This creates a distinction between

$$\Omega_i^{\text{formal}}(t) \qquad \text{and} \qquad \Omega_i^{\text{substantive}}(t). \tag{73}$$

The distinction matters for near-critical response because trajectory closure can operate through indirect channels. A group can retain formal rights while losing practical access to institutions, information, mobility, finance, or collective organization needed to exercise those rights. Conversely, an institutional reform can expand substantive options even if the formal menu of choices changes little.

Working Claim 15.3 (Conversion-sensitive option analysis)**.** Near-critical option analysis should examine how transitions modify the conditions under which actors can convert formal possibilities into practicable courses of action. Counting nominal alternatives can conceal substantial contractions or expansions of agency.

This claim also places a limit on aggregate option measures. A collective system can appear to possess many future pathways while some actors are unable to enter most of them. The distribution of conversion capacity therefore belongs to the justice profile rather than being treated as a purely technical feature of adaptability.

15.3 Agency, secure options, and structural dependence

Practical availability remains incomplete if an option exists only at another actor's uncontrolled discretion. Republican accounts of freedom as non-domination emphasize this structural distinction: a person can experience little actual interference while remaining vulnerable to a power that may interfere arbitrarily (Lovett 2026). The present paper uses this idea comparatively to distinguish a temporarily available option from a politically secure option.

Definition 15.4 (Secure political option)**.** A *secure political option* is a substantively reachable option whose continued availability does not depend primarily on the uncontrolled discretion of another actor and whose relevant institutional conditions possess a reasonable degree of procedural durability.

Security matters especially near criticality because small institutional changes can transfer control over later pathways. A transition can expand the number of options controlled by a central authority while contracting the options independently available to affected groups. Aggregate adaptability can therefore coexist with increased dependence.

Agency adds another dimension. An actor can possess several options yet have little role in deciding which collective trajectory becomes binding. The paper therefore distinguishes *having options* from *participating in trajectory formation*. This distinction connects with political legitimacy: questions about who may exercise coercive political power and under which conditions cannot be answered by dynamical efficacy alone (Peter 2023).

Definition 15.5 (Trajectory agency)**.** *Trajectory agency* is the practically effective capacity of an affected actor or group to participate in the interpretation, selection, contestation, and later revision of collective pathways that materially structure its future option space.

Trajectory agency is relational. It depends on institutional access, information, timing, representation, coordination capacity, and the distribution of agenda-setting power. Near criticality, timing can be especially important: participation offered after a transition has become strongly self-reinforcing can be formally present while having little capacity to reopen the pathway.

Proposition 15.6 (Option–agency separation)**.** *A transition can preserve several formal or substantive options while reducing trajectory agency if affected actors lose meaningful capacity to shape which options remain open, which become binding, or how later revision can occur.*

This proposition adds a procedural dimension to option preservation. A political system cannot be assessed solely by the number of futures still reachable from the perspective of an external analyst. The distribution of standing over those futures also matters.

15.4 Revisability after collective commitment

Revisability names the capacity to alter a chosen trajectory after collective commitment. The term should be distinguished from perpetual indecision. Political coordination often requires commitments, investments, rules, and expectations that acquire some durability. A system in which every decision can be instantly reversed can become unable to coordinate long-horizon action. The problem near criticality is therefore not commitment itself. It is the relation between commitment, learning, and the practical possibility of correction.

Definition 15.7 (Political revisability)**.** *Political revisability* is the capacity of a political system and its affected actors to review, alter, suspend, compensate, or reopen a collective trajectory when relevant evidence, values, or conditions change, without requiring prohibitive social or political cost.

The definition has both dynamical and institutional components. Dynamically, hysteresis can make return difficult even when actors agree that revision is desirable. Institutionally, legal rules, sunk investments, organizational dependence, and distributions of power can raise the cost of reversal. These components can reinforce one another.

Adaptive policy-pathway research provides a useful applied comparison because it explicitly treats plans as sequences that can change over time and uses pathways, signposts, triggers, and contingency actions to represent future adaptation under deep uncertainty (Haasnoot et al. 2013). The present framework extends the normative question beyond strategic flexibility. It asks who retains the capacity to initiate revision, whose interests are represented in the trigger conditions, and who bears the costs of keeping alternatives open.

A simple conceptual quantity is the *revision cost*

$$\mathcal{C}_i^{\text{rev}}(\gamma, t) = \mathcal{C}(\text{reopening or altering trajectory } \gamma \text{ for actor } i \text{ at time } t)\,. \tag{74}$$

Trajectory closure corresponds in part to a growth of $\mathcal{C}_i^{\text{rev}}$ for relevant actors and alternatives. The costs need not be monetary. They can include institutional displacement, collective action requirements, legal barriers, informational loss, reputational exposure, or the disappearance of coordinating organizations.

Proposition 15.8 (Revisability value under uncertainty)**.** *Where uncertainty about long-run consequences remains substantial and later correction is plausibly valuable, lower revision costs can have independent political significance. This significance remains conditional on the burdens created by keeping alternatives open and does not imply that every commitment should remain easily reversible.*

The qualification is important. Section 13 already showed that delay and flexibility can impose costs on actors currently exposed to harm. Revisability is therefore a value to be balanced within a broader justice and risk profile, rather than a universal instruction to avoid commitment.

15.5 Asymmetric option contraction

Collective option space can expand while the option space of particular actors contracts. This is one reason to avoid treating $|\Omega_{\text{collective}}|$ as a sufficient indicator of political openness. A system can become more adaptable from an institutional center while making peripheral actors more dependent, or preserve many aggregate pathways by concentrating the burdens of flexibility on a smaller population.

Let

$$\Delta\Omega_i = \Omega_i(t+1) - \Omega_i(t) \tag{75}$$

be interpreted qualitatively as a change in actor i's practical option structure. Near-critical justice requires attention to the distribution of such changes across actors rather than only to their aggregate sum. The notation is schematic because option spaces are structured sets whose value cannot normally be represented by simple subtraction.

Definition 15.9 (Asymmetric option contraction)**.** *Asymmetric option contraction* occurs when a transition reduces the substantive, secure, or revisable pathways available to some affected actors more strongly than to others, especially where the actors bearing the contraction have limited standing in the decision that produced it.

The temporal dimension extends this concern across generations. Work on intergenerational justice emphasizes that present actions shape the conditions and opportunities available to future persons (Meyer 2021). Near-critical transitions sharpen this problem when current actors exploit a window of leverage to establish persistent institutional or ecological conditions whose revision costs are transferred forward in time.

The reverse pattern is also possible. Preserving every present option can impose large future burdens when delay allows an accumulating hazard to deepen. The relevant normative question is therefore distributional and temporal: whose options are preserved, whose are consumed, for how long, and under what conditions of voice and compensation.

Research Obligation 15.10 (Distributed option-space record)**.** Where a near-critical intervention is expected to change long-run reachability, analysis should record, where feasible, the actors whose options are likely to expand or contract, the security and substantive availability of those options, the anticipated revision costs, and the temporal location of the associated burdens. Aggregate flexibility should be reported separately from its distribution.

This obligation extends the transition-justice record introduced in Section 12. It also supplies a bridge from normative analysis to institutional design: an institution that claims to preserve flexibility needs procedures for identifying whose flexibility is actually being preserved.

15.6 Qualified option preservation

The preceding analysis supports a qualified orientation toward option preservation. The orientation is deliberately weaker than an instruction to maximize the number of available futures. Options differ in substantive value, security, distribution, feasibility, and the harms involved in keeping them open. Some pathways depend on domination or severe ongoing harm. Some nominal options consume resources needed to protect more basic forms of agency. Some commitments are necessary to create capabilities that would otherwise remain unavailable.

The option-preservation orientation can therefore be stated as a conditional constraint:

Proposition 15.11 (Qualified option-preservation orientation)**.** *Under substantial uncertainty about long-run trajectories, and where the costs of continued openness remain ethically acceptable, near-critical response has reason to avoid unnecessary and highly asymmetric closure of substantive, secure, and revisable options. The strength of this reason decreases when delay or continued openness imposes severe ongoing harm, entrenches domination, or consumes the agency of affected actors.*

This proposition avoids equating openness with goodness. It also avoids the opposite assumption that a coherent and stable trajectory should be preferred because closure simplifies coordination. The appropriate question concerns the quality and distribution of the possibilities that remain, together with the costs of preserving them.

Remark 15.12 (Option quality and option quantity)**.** A larger option set can coexist with weaker agency when the added options are nominal, insecure, unusable, or controlled by others. Near-critical analysis should therefore distinguish the quantity of alternatives from their substantive quality and political distribution.

This point connects directly to the aesthetic analysis of Section 14. Open configurations can look less coherent than consolidated ones, yet incompleteness can preserve pathways for later learning and participation. It also connects to *wuwei*: a non-coercive response can sometimes protect the conditions of later agency by resisting premature closure, while other situations can require active protection because continued waiting steadily destroys those same conditions.

15.7 Agency-preserving response and institutional transition

Option space, agency, and revisability can now be integrated into the larger near-critical response framework. A response is *agency-preserving* to the extent that it protects meaningful participation in later trajectory formation, maintains practicable routes of revision where feasible, and avoids shifting large option losses onto actors with little control over the transition.

Definition 15.13 (Agency-preserving response)**.** An *agency-preserving response* is a mode of near-critical participation that gives explicit weight to the substantive options, secure standing, revision capacity, and future participation of affected actors while addressing the risks that make response necessary.

Agency preservation is therefore compatible with intervention, buffering, restraint, or temporary commitment depending on the configuration. It is not a synonym for minimal state action or maximal decentralization. A protective intervention can expand agency when it prevents an irreversible loss of basic capabilities. A restrained response can preserve agency when forceful steering would prematurely lock several groups into a trajectory they cannot later contest. The assessment remains relational and case-sensitive.

The section's contribution can be summarized by separating four questions:

(i) which future pathways remain practically reachable;

(ii) which actors can substantively use and securely retain those pathways;

(iii) which actors possess effective standing in trajectory formation; and

(iv) how costly later revision becomes after collective commitment.

These questions complement the earlier analyses of dynamics, epistemics, justice, ethics, and aesthetics. Together they shift attention from the selection of a terminal state toward the political conditions under which collective becoming remains contestable and revisable.

The next section turns from these normative and analytical requirements to institutional architecture. If near-critical response is to preserve sensing, response diversity, agency, and revisability while limiting catastrophic propagation, then institutions require structures capable of distributing observation, buffering cascades, coordinating heterogeneous responses, and reopening decisions when the conditions supporting them change.

16 Institutional Architectures for Near-Critical Response

The preceding sections developed near-critical response as a problem of susceptibility, perturbation–response relations, trajectory closure, justice, irreversibility, aesthetic judgment, option space, and agency. This section translates those analytical requirements into an institutional register. Its structural role is limited and practical: it does not specify an ideal constitutional order, and it does not assume that one organizational form is universally suited to near-critical conditions. The objective is to identify architectural properties that can help political systems sense changing susceptibility, contain destructive cascades, preserve heterogeneous response capacity, keep important decisions revisable, and retain the standing of affected actors while consequences are still unfolding.

The discussion proceeds through seven institutional functions: distributed sensing, buffering and modularity, differentiated response speed, reversible action and review, polycentric coordination, participation and transition standing, and institutional memory. These functions are evaluated through the response-centered framework developed earlier rather than through a single criterion of efficiency. Adaptive-governance research provides useful comparative resources because it emphasizes multi-level organization, learning, networks, and reorganization under abrupt change (Folke et al. 2005). Polycentric-governance scholarship similarly directs attention to multiple centers of decision making and mutual adjustment (Ostrom 2010). The present argument adds a more specific question: how should institutional architecture behave when susceptibility is elevated and policy action can itself become an amplified perturbation?

Table 6 summarizes a central point. Near-critical institutional capacity is distributed across functions that can pull in different directions. Modularity can reduce cascade risk while weakening common coordination. Speed can protect a boundary condition while shrinking participation. Revisability can preserve correction capacity while undermining commitment. The relevant design problem therefore concerns combinations and sequencing rather than maximization of any single property.

16.1 Response architecture and institutional role

An institution does more than select a policy instrument. It structures what is observed, which perturbations become legible, whose interpretations enter the decision, which response modes are available, how quickly action can occur, and how easily the resulting decision can later be reopened. Near criticality these features become dynamically significant because the institution helps shape the response field through which a perturbation travels.

Definition 16.1 (Near-critical response architecture)**.** A *near-critical response architecture* is the institutional configuration through which perturbations are observed, interpreted, received, transformed, and reviewed under conditions of elevated susceptibility. It includes the distribution of sensing, authority, buffers, response channels, revision procedures, participation, and institutional memory.

This definition treats architecture as an arrangement of relations rather than a list of organizations. Two formally different states can possess similar response architectures when they rely on the same data pipelines, professional models, escalation routines, and centralized decision channels. Conversely, one formal organization can contain internally heterogeneous response units with distinct information sources and revision procedures.

A schematic representation is

$$\mathfrak{A}_t = (\mathcal{S}_t, \mathcal{B}_t, \mathcal{L}_t, \mathcal{R}_t, \mathcal{P}_t, \mathcal{M}_t)\,, \tag{76}$$

where $\mathcal{S}_t$ denotes sensing arrangements, $\mathcal{B}_t$ buffering and boundary structures, $\mathcal{L}_t$ latency and sequencing rules, $\mathcal{R}_t$ revision and reopening capacity, $\mathcal{P}_t$ participation and standing, and $\mathcal{M}_t$ institutional memory. Equation (76) is an analytical inventory. It does not imply that these dimensions can be reduced to a single institutional score.

Table 6: Institutional functions for near-critical response. The entries identify architectural possibilities and characteristic tensions rather than a universal design template.

Institutional function	**Illustrative architectural form**	**Near-critical purpose**	**Characteristic tension**
Distributed sensing	Independent observatories, local reporting channels, cross-domain monitoring	Reduce dependence on one model and reveal heterogeneous response traces	Fragmented evidence can obstruct synthesis or be selectively ignored
Buffering and modularity	Firebreaks, reserve capacity, compartmentalized channels, redundancy	Limit propagation and preserve local functionality during disturbance	Excessive insulation can block coordination and conceal transferred burdens
Differentiated response speed	Fast protective mechanisms paired with slower trajectory-changing procedures	Match decision latency to the persistence and uncertainty of the hazard	Emergency speed can normalize exceptional authority or bypass affected actors
Reversibility and review	Sunset clauses, staged authorization, policy triggers, reopening procedures	Preserve correction capacity when models or conditions change	Frequent reopening can weaken credible commitments and create strategic instability
Polycentric coordination	Multiple decision centers linked by communication and escalation protocols	Reduce correlated error while retaining capacity for coordinated action	Overlap can create ambiguity, veto congestion, or unequal accountability
Participation and transition standing	Rapid consultation, representative panels, affected-group review, contestation channels	Keep trajectory formation connected to actors bearing risk and option loss	Time pressure can narrow participation or privilege already organized actors
Institutional memory	Versioned decisions, rationale records, trigger histories, response audits	Support learning across repeated perturbations and enable later accountability	Records can become performative signals, bureaucratic burdens, or tools of strategic blame

Proposition 16.2 (Architectural mediation)**.** *The political effect of a perturbation near criticality is partly mediated by the institutional architecture through which it is sensed, transmitted, interpreted, and answered. Institutional design therefore changes the perturbation–response relation even when the initiating disturbance remains the same.*

The proposition connects institutional analysis to Section 4. A disturbance that propagates through a single synchronized channel can produce a different trajectory from the same disturbance entering a system with buffers, independent observation, and several response modes. The claim remains conditional: architecture shapes the field without fully determining the result.

16.2 Distributed sensing and epistemic independence

Section 6 described susceptibility as a latent and direction-dependent property inferred from noisy traces. Institutional architecture should therefore preserve some independence among the channels through which those traces are generated and interpreted. A system containing many reporting units can still be epistemically centralized when all units use the same indicators, forecasting model, incentive structure, or upstream data source.

Let the sensing system contain m channels

$$\mathcal{S}_t = \{s_t^{(1)}, s_t^{(2)}, \ldots, s_t^{(m)}\}. \tag{77}$$

The relevant property is not simply m. It is the degree to which errors, blind spots, and interpretive assumptions are imperfectly correlated. Local knowledge, statistical indicators, administrative data, professional judgment, and reports from affected communities can therefore have architectural value when they expose different aspects of the same evolving configuration. Adaptive-governance research has similarly emphasized multiple knowledge systems and cross-level learning (Folke et al. 2005).

Working Claim 16.3 (Epistemic independence buffer)**.** Where susceptibility is uncertain and evidence channels share common failure modes, preserving partially independent sensing and interpretation can reduce the probability that one mistaken classification becomes a system-wide response trigger. The value of such independence depends on institutions that can compare, contest, and synthesize heterogeneous evidence without erasing its source-specific uncertainty.

Independence alone is insufficient. Fragmented observation can produce parallel accounts that never enter a common decision process. Near-critical architecture therefore requires both differentiation and translation. Bridging units can assemble conflicting signals, preserve disagreement in the record, and identify which response conclusions are robust across interpretations. Their task is closer to epistemic coordination than to declaring one definitive model of the system.

This point also places a constraint on dashboards and composite risk indices. Compression is often necessary for decision making, yet a single scalar alert can conceal the structure of uncertainty that matters most near criticality. Institutional summaries should therefore retain access to the underlying signals, model assumptions, and disagreement that generated the summary.

16.3 Buffers, modularity, and institutional firebreaks

Near-critical systems can be endangered by rapid propagation across tightly coupled subsystems. Network resilience research identifies redundancy, diversity, and modularity as structural properties that can limit the spread of disturbance under some conditions (Kharrazi et al. 2020). The political analogue is an institutional architecture in which a failure or misclassification in one channel does not automatically force all connected units into the same response.

Definition 16.4 (Institutional firebreak)**.** An *institutional firebreak* is a boundary, buffer, delay, reserve, or procedural separation designed to reduce the probability that a local perturbation, policy error, or response cascade immediately propagates across the wider political system.

Examples can include temporary separation of decision channels, reserve capacity that prevents immediate dependence on one provider, procedural delays between signal detection and system-wide escalation, or jurisdictional boundaries that allow local experimentation without instant universalization. The category is functional rather than organizational. A legal rule, reserve fund, communications protocol, technical interface, or institutional division can each operate as a firebreak.

The water metaphor from Section 10 is useful only at a limited level here. A riverbed contains banks, channels, pools, and roughness that redistribute flow without specifying each movement of water. Institutional firebreaks can similarly shape transmission without selecting a complete final trajectory. The analogy stops where political burdens and agency begin. A barrier that protects one subsystem can divert pressure toward another, and that displacement has distributive significance.

Proposition 16.5 (Cascade containment with burden tracing)**.** *Architectural measures that reduce cascade propagation should be evaluated jointly with the burdens they redirect. A firebreak can improve system-level containment while worsening exposure for actors located beyond the protected boundary.*

This proposition links resilience to the transition-justice analysis of Section 12. Modularity and redundancy can be valuable, while their distributional incidence remains open to scrutiny. Institutional architecture therefore needs burden tracing alongside cascade analysis.

16.4 Differentiated latency and two-speed response

Decision speed becomes unusually consequential near criticality. Excessive latency can allow a destructive feedback loop to deepen, while excessive speed can convert uncertain evidence into premature trajectory closure. A single institutional tempo is therefore poorly suited to a field containing both urgent boundary threats and uncertain long-horizon transitions.

Let τ_a denote the latency of response mode a. Near-critical design requires a relation between response latency and the persistence profile of the hazard rather than universal minimization of τ_a. A protective response to an imminent irreversible harm can require short latency, while a response that selects a long-run institutional trajectory can require additional time for verification, participation, and review.

Definition 16.6 (Differentiated latency)**.** *Differentiated latency* is the institutional capacity to assign different response times to different classes of action according to their persistence, reversibility, uncertainty, and distributive reach.

This suggests a provisional two-speed architecture. The first layer contains fast, bounded mechanisms that protect minimum boundary conditions or interrupt an evidently dangerous cascade. The second contains slower procedures for trajectory-changing decisions whose consequences are durable, widely distributed, or difficult to revise. The layers should remain connected: rapid protection creates time for slower judgment, and slower review can terminate, revise, or legitimate the temporary measures adopted under urgency.

Working Claim 16.7 (Latency differentiation)**.** Near criticality, institutional speed has value when it is differentiated by the type of consequence at stake. Rapid protective action and slower trajectory-forming action can coexist within one architecture when the former is bounded, reviewable, and prevented from silently becoming the latter.

The final clause is essential. Emergency mechanisms often create their own institutional path dependence. Temporary authority, data infrastructures, security procedures, or exceptional restrictions can become difficult to remove after the initiating disturbance has passed. A two-speed architecture therefore requires procedures that distinguish temporary protection from durable constitutional or distributive change.

16.5 Reversibility, triggers, and reopening procedures

Section 15 treated revisability as a politically significant property under uncertainty. Institutions can support revisability through staged authorization, sunset provisions, review dates, contingency plans, and predefined routes for reopening decisions. Dynamic Adaptive Policy Pathways offers a useful planning comparison because it uses signposts, triggers, alternative pathways, and adaptation over time under deep uncertainty (Haasnoot et al. 2013). Near criticality, however, trigger design needs an additional caution: a trigger based on an uncertain indicator can itself become part of a self-reinforcing response cascade.

A trigger should therefore be separated conceptually from the action it initiates. Let

$$z_t \in \mathcal{Z} \quad \Longrightarrow \quad \mathcal{Q}_t, \tag{78}$$

where z_t is an observed signpost and $\mathcal{Q}_t$ is a review or bounded response procedure. Equation (78) deliberately avoids an automatic mapping from a warning signal to a strong intervention. Some triggers can activate further observation, broaden consultation, increase reserves, or prepare alternative responses before they authorize a trajectory-changing measure.

Proposition 16.8 (Trigger discipline)**.** *Where warning indicators are noisy, performative, or capable of generating correlated responses, institutional triggers should be calibrated to the persistence and*

reversibility of the action they activate. Evidence sufficient to begin review can remain insufficient to authorize durable trajectory closure.

Reopening procedures matter after action as well. A response should record the conditions under which it remains justified and the observations that would support modification or withdrawal. This converts reversibility from an abstract preference into an institutional practice. It also limits hindsight bias by preserving the reasons and uncertainty present at the time of decision.

16.6 Polycentric coordination and response correlation

The response-repertoire analysis of Section 9 identified correlated response as a source of vulnerability. Polycentric arrangements can provide one institutional means of maintaining partially independent responses while allowing coordination across centers (Ostrom 2010). The relevant property is neither maximal decentralization nor maximal centralization. It is the architecture of coordination among centers that possess differentiated information, authority, and response capabilities.

A highly centralized architecture can synthesize information and act quickly, yet it can also amplify one model error across the entire system. A highly fragmented architecture can preserve experimentation and local adaptation, yet it can fail to coordinate around cross-boundary risks. Near-critical response therefore requires attention to the correlation structure of institutions. Several agencies with different names can remain functionally synchronized when they share the same incentives and information. Several jurisdictions can remain productively differentiated when they exchange information while retaining meaningful local discretion.

Definition 16.9 (Coordinated heterogeneity)**.** *Coordinated heterogeneity* is an institutional condition in which multiple response centers retain substantively distinct information or action capacities while maintaining communication, escalation, and mutual-learning channels sufficient for problems that cross their boundaries.

This concept sharpens the earlier portfolio analogy. The aim is to reduce correlated failure without turning heterogeneity into institutional isolation. Adaptive-governance scholarship emphasizes cross-scale connections and bridging organizations precisely because decentralized adaptation and broader coordination need mechanisms of connection (Folke et al. 2005). Complex-governance work likewise cautions that governance capacity under complexity depends on interactions among institutions, organizations, and actors rather than one steering center (Duit and Galaz 2008; May 2022).

16.7 Transition standing and institutional participation

Section 12 introduced transition standing as the claim of affected actors to participate in or contest decisions that can alter their long-run political possibilities. Institutional architecture gives that claim procedural form. Near-critical conditions create a genuine tension: long deliberation can be costly when harms are accelerating, while compressed decision processes can concentrate trajectory selection in actors whose own exposure is limited.

Participation therefore needs several temporal forms. Some processes can occur before a crisis through standing representative bodies, pre-agreed procedural constraints, and public specification of emergency authorities. Some can occur during the critical interval through rapid consultation, affected-group review, expert disagreement panels, and contestation channels. Other forms occur after a provisional response through mandatory review, compensation claims, and reopening rights.

Working Claim 16.10 (Temporal distribution of participation)**.** Where near-critical conditions constrain the time available for deliberation, legitimate participation can be distributed across preparation, immediate response, and subsequent review. Compressed decision time does not erase the need for transition standing; it changes the institutional locations through which that standing can be exercised.

This formulation also prevents participation from becoming a ceremonial delay. The relevant question is whether affected actors possess meaningful routes to shape, contest, or reopen consequential decisions. Institutions should record whose participation was absent, whose burdens increased, and which remedies remain available after the immediate critical interval.

16.8 Institutional memory and reopening capacity

Near-critical response is iterative. Each perturbation, intervention, and revision can provide information about the response field. That information is lost when institutions preserve only the final decision. Institutional memory should therefore include the observations, alternative interpretations, uncertainty judgments, response sequence, distributive expectations, and review conditions that accompanied the decision.

Let a decision record be represented schematically as

$$\mathcal{H}_t^{\text{inst}} = (E_t, U_t, A_t, J_t, O_t, C_t)\,, \tag{79}$$

where E_t denotes the evidence available at the time, U_t uncertainty, A_t considered actions, J_t justice and legitimacy considerations, O_t anticipated option-space effects, and C_t conditions for continuation, revision, or withdrawal. The notation is a record schema rather than a demand for complete documentation.

Institutional memory serves three purposes. First, it supports learning by allowing later observers to compare expected and realized response dynamics. Second, it protects revisability by preserving the conditions under which a measure was adopted. Third, it supports accountability by making visible how risk and uncertainty were distributed at the time of decision. These purposes can conflict with speed, confidentiality, and strategic flexibility, so the recording architecture itself requires proportionality.

Definition 16.11 (Institutional reopening capacity)**.** *Institutional reopening capacity* is the practical ability of a political system to revisit a near-critical response when its evidential basis, dynamical context, distributive effects, or legitimacy conditions materially change.

Reopening capacity differs from routine policy turnover. It requires that the institution retain pathways, authority, information, and procedural standing sufficient to reconsider a decision whose own effects may already be creating institutional hysteresis. Sunset clauses without a functioning review body, records without decision authority, or formal appeals without substantive access provide weak reopening capacity.

16.9 Architectural orientation and transition to evidence

The institutional discussion can now be condensed into a provisional orientation. Near-critical architectures should seek combinations of epistemic plurality, cascade containment, differentiated response speed, revisability, coordinated heterogeneity, transition standing, and institutional memory. These properties remain conditional. A system facing immediate catastrophic harm can legitimately compress some procedures; a system facing deep uncertainty about a highly persistent transition can require stronger barriers against rapid trajectory closure.

Proposition 16.12 (Near-critical institutional orientation)**.** *An institution is better positioned to respond near criticality when it can observe through partially independent channels, contain local failures without simply exporting them, vary response speed according to consequence, preserve credible routes of revision, coordinate heterogeneous centers, maintain the standing of affected actors, and learn from the historical record of its own responses. The proposition specifies an orientation for inquiry rather than a complete institutional ranking.*

The section therefore resists the conversion of near-criticality into a new technocratic design doctrine. Institutional architecture can reduce some forms of correlated error and irreversible closure while creating other forms of power, delay, or exclusion. The purpose of architecture is to shape the conditions of response while keeping its own limits visible.

The next section turns to evidence and policy. Once institutions are understood as part of the perturbation field, evidence cannot be treated as a neutral input that mechanically selects action. Near criticality, locally valid evidence can remain epistemically informative while becoming dynamically unsafe or normatively insufficient as a basis for intervention. Section 17 develops this regime-sensitive relation between evidence and action.

17 Evidence, Policy, and Regime-Sensitive Action

This section develops a regime-sensitive account of evidence-to-action relations near criticality. Its role is to identify a limit that remains even when the available evidence is methodologically strong, locally relevant, and sincerely used. The problem concerns neither the rejection of evidence nor the familiar observation that policy is shaped by values and politics. It concerns a more specific possibility: evidence that is informative about a present response relation can become an insufficient basis for intervention when the intervention itself can alter the regime in which that relation holds. The section has five objectives. First, it distinguishes evidential validity from intervention safety. Second, it examines the transport of historical evidence into near-critical regimes. Third, it treats policy action and policy communication as potentially performative components of the response field. Fourth, it formulates a regime-sensitive sequence for evidence use in which dynamical and normative questions remain separate from evidential quality. Fifth, it develops the possibility that improved evidence about leverage can sometimes support restraint, buffering, or further observation rather than stronger optimization.

Evidence-based policy scholarship already resists a mechanical relation between research findings and public action. Head argues that policy decisions emerge through the interaction of research, practical experience, political judgment, and contested values rather than deduction from a single evidence base (Head 2008; Head 2010). Cairney similarly emphasizes bounded attention, ambiguity, and the complex institutional environments in which evidence is interpreted and used (Cairney 2016a; Cairney 2016b). Parkhurst develops the related idea of the good governance of evidence, in which questions of evidence quality remain distinct from the political and procedural conditions of its use (Parkhurst 2017). Cartwright and Hardie add a transportability problem: evidence that an intervention worked elsewhere is insufficient without knowledge of the support factors that allow the causal relation to operate in the target setting (Cartwright and Hardie 2012). The argument developed here builds on these limitations while adding a dynamical one. Near criticality, the target setting can change partly because of the attempted use of the evidence itself.

Figure 10 summarizes the problem. Evidence can be reliable for one inferential task while remaining incomplete for the decision task created by near-critical susceptibility. The issue is therefore not a simple hierarchy between "good evidence" and "bad evidence." It is the relation among evidence, regime, intervention, and the possibility that this relation changes during use.

17.1 Evidential validity and intervention safety

Let E_t denote an evidential basis available at time t, and let $\widehat{R}_{r_t}(u)$ denote an estimated response to intervention u under regime r_t. A conventional local policy inference can be represented as

$$E_t \Longrightarrow \widehat{R}_{r_t}(u) \Longrightarrow \text{policy judgment about } u. \tag{80}$$

The first implication can be methodologically strong. The evidence may support a credible estimate of direction, short-run effect, or conditional response. The second implication is nevertheless a different operation. It concerns whether acting on that relation is dynamically safe, normatively warranted, and institutionally legitimate.

Sections 7 and 12 already separated local epistemic validity from reversibility and normative warrant. The same distinction can be expressed for evidence-based policy through three questions:

(i) Is the evidence sufficiently reliable for the stated inferential task?

(ii) Is the intervention likely to preserve a response regime in which that inference remains relevant over the decision horizon?

(iii) Is the use of the intervention justified given its distributional, procedural, option-space, and legitimacy effects?

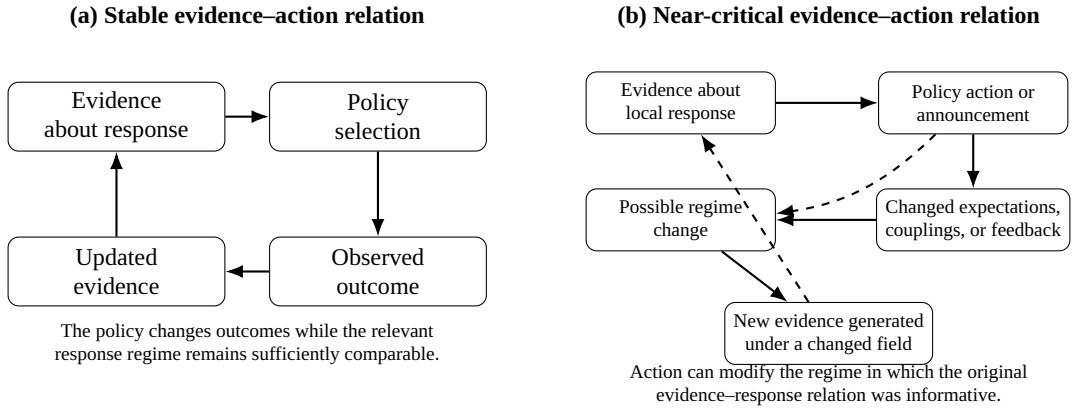

Figure 10: Two evidence–action structures. Panel (a) represents a setting in which policy changes outcomes while the response regime remains sufficiently comparable for iterative updating. Panel (b) represents a near-critical possibility in which action or communication changes expectations, coupling, feedback, or attractor accessibility, so subsequent evidence is generated in a field that can differ materially from the one in which the initial evidence was obtained. The figure is conceptual and does not imply that every near-critical policy intervention causes a regime transition.

A positive answer to the first question does not settle the other two.

Definition 17.1 (Critical-regime instability of evidence–action relations)**.** *Critical-regime instability of evidence–action relations* is the condition in which evidence remains informative about a local or current response relation while policy action based on that relation can materially change the feedback structure, attractor accessibility, coupling pattern, or regime in which the evidential relation was estimated.

This definition identifies a limitation of evidence-to-action inference rather than a limitation of empirical inquiry itself. Evidence retains value. Indeed, near criticality can make high-quality observation more important. The problem is that the informational value of evidence and the action warrant inferred from it can move in different directions.

Proposition 17.2 (Local validity and action safety)**.** *Near a plausible critical transition, locally valid evidence can remain informative while the set of interventions for which that evidence provides an adequate basis for safe and legitimate action contracts. Improved local estimation can therefore coexist with reduced action admissibility.*

The proposition is conditional. If the intervention is weakly coupled to the transition mechanism, highly reversible, normatively authorized, or protected by reliable buffers, the action region can remain broad. The important point is that evidential quality alone does not determine that region.

17.2 Historical evidence and regime transport

Evidence-based policy commonly relies on some form of transport: a relation observed in one population, time, program, or institutional context is used to inform a decision elsewhere. Cartwright and Hardie emphasize that such transport requires attention to the support factors that allow a policy to work in the target context (Cartwright and Hardie 2012). Near criticality adds a temporal form of the same problem. The relevant target context may be the same jurisdiction or institution, but its dynamical regime can have changed.

Let $\mathcal{D}_{\text{hist}}$ denote historical observations generated under regime r_0, and let the current policy decision occur under a configuration r_t. Evidence transport can be written schematically as

$$\mathcal{T}_E : \mathcal{D}_{\text{hist}}(r_0) \mapsto \widehat{R}_{r_t}(u), \tag{81}$$

where $\mathcal{T}_E$ represents the assumptions that make historical evidence relevant to the present decision. Near criticality, these assumptions require special attention because response elasticities, feedback signs, propagation paths, and recovery rates can change with the regime.

A policy can therefore be supported by excellent historical evidence and still face a transport problem. The issue is neither that the original study was wrong nor that the past is irrelevant. The issue is that a previously stable conditional relation can enter a configuration in which the same intervention has a different system-level meaning. Evidence about ordinary marginal effects can be weak evidence about basin crossing. Evidence about average response can be weak evidence about tail propagation. Evidence about a previous recovery rate can be weak evidence about recovery after a feedback architecture has changed.

Working Claim 17.3 (Regime-sensitive transport)**.** The transportability of policy evidence depends partly on whether the causal and response conditions relevant to the evidence remain sufficiently stable across the source and target regimes. Proximity to a transition is therefore a potential transport condition rather than merely an additional contextual variable.

This claim extends the contextualism already present in evidence-based policy analysis. It also explains why repeated success can sometimes provide less security than it appears to provide. If the system has been moving toward a threshold while the policy was repeatedly effective, the historical record may mainly describe the pre-transition regime. An apparent sequence of successful small adjustments does not establish that the next adjustment remains in the same response domain.

17.3 Performative evidence and reflexive response

Political evidence is often communicated. Forecasts, risk assessments, public warnings, policy commitments, and official classifications can change the expectations of actors who are themselves components of the system. Under ordinary conditions this reflexivity can be part of implementation. Near criticality it can acquire additional significance because expectation shifts can change timing, coordination, withdrawal, investment, alliance formation, or other forms of coupling.

Let $m(E_t)$ denote a public or semi-public representation of evidence. A reflexive political field can then be represented as

$$x_{t+1} = F(x_t, u_t, m(E_t), \mathbb{E}_t[x_{t+1}], G_t) + \eta_t, \tag{82}$$

where $\mathbb{E}_t[x_{t+1}]$ represents actors' expectations and G_t the relevant interaction structure. Equation (82) is a formal sketch. Its purpose is to keep visible that evidence can affect the system through interpretation and anticipation in addition to any direct policy instrument.

Definition 17.4 (Evidence performativity near criticality)**.** *Evidence performativity near criticality* is the condition in which the production, communication, or institutional use of evidence changes actor expectations or relations in ways that materially alter the response field to which the evidence refers.

Performativity can move in several directions. A warning can coordinate protective action and reduce transition risk. It can also accelerate a run, induce pre-emptive escalation, synchronize previously independent actors, or change the strategic value of waiting. A public declaration that one trajectory is "likely" can become a focal point around which behavior reorganizes. These possibilities strengthen the case for distinguishing observation from the institutional effects of observation. They also connect Section 17 with Section 8: evidence use is itself one possible perturbation channel.

The implication is not secrecy. Suppressing evidence can generate its own risks, legitimacy problems, and information asymmetries. The implication is that communication design belongs to the response analysis. Who receives the evidence, when, in what form, with what uncertainty representation, and under which institutional commitments can matter to the resulting dynamics.

Table 7: Regime-sensitive dimensions of evidence-to-action assessment. The entries identify distinct questions rather than a ranking or scoring system.

Dimension	Primary question	Near-critical complication
Evidential validity	What does the available evidence support for the stated inferential task?	Strong local inference can coexist with weak long-horizon prediction.
Regime transport	Do the causal and response conditions of the evidence remain relevant in the present configuration?	Proximity to transition can change feedback, propagation, and recovery relations.
Intervention sensitivity	How strongly can the proposed action couple to susceptible directions?	Small actions can have disproportionate or persistent trajectory effects.
Performativity	How can evidence communication and policy anticipation alter behavior?	The evidence-use process can become part of the perturbation field.
Reversibility	What credible routes remain for correction, withdrawal, or reopening?	A modest initial action can cross into a persistent response regime.
Normative warrant	Who bears risk, loses options, participates, and possesses standing?	Low instrumental cost can coexist with high distributive or legitimacy consequences.

17.4 Regime-sensitive evidence-to-action assessment

The previous arguments suggest an evidence-to-action assessment that is broader than conventional questions of methodological quality. The sequence is not a mechanical checklist and does not assign a universal ordering to policy goals. It makes several distinct judgments visible before they are compressed into a single recommendation.

Table 7 does not diminish the role of evidence. It makes evidence more specific. Instead of asking only whether a policy is "evidence based," inquiry asks what the evidence establishes, which regime it speaks to, how action changes the response field, and which normative judgments remain external to the evidence itself. This orientation is compatible with Head's account of multiple evidentiary lenses and with Parkhurst's distinction between evidence quality and the governance of evidence (Head 2008; Parkhurst 2017). It also retains Cairney's emphasis on ambiguity and bounded policy attention while adding regime instability as a separate source of action difficulty (Cairney 2016b).

Definition 17.5 (Regime-sensitive action)**.** *Regime-sensitive action* is intentional political response that conditions its use of evidence on the estimated response regime, transition uncertainty, intervention-induced changes in the field, reversibility, and the normative standing of affected actors. It can support intervention, buffering, staged experimentation, observation, delay, or restraint according to the configuration.

The definition avoids equating regime sensitivity with caution in every case. A rapidly escalating irreversible harm can justify fast protective action. Conversely, a highly leveraged but normatively contested intervention can justify a higher threshold for deliberate steering. The orientation is conditional because near-critical response remains a relation among the system, the perturbation, the available evidence, and the actors who bear consequences.

17.5 Evidence, leverage, and restraint

A common intuition in policy analysis is that improved evidence should reduce uncertainty and thereby strengthen the case for action. This relation often holds. Near criticality, however, one kind of improved

evidence can reveal that an intervention is more consequential than previously assumed. Better knowledge of susceptibility can therefore increase both efficacy expectations and the ethical burden of use.

Let $L_t(u)$ denote estimated trajectory leverage and let $W_t(u)$ denote the strength of normative and institutional warrant for using u. A regime-sensitive decision cannot identify $L_t(u)$ with $W_t(u)$. In particular,

$$L_t(u) \uparrow \qquad \text{does not entail} \qquad W_t(u) \uparrow . \tag{83}$$

The separation is especially important when leverage derives from fragility, information asymmetry, concentrated exposure, or a narrow transition window.

Proposition 17.6 (Evidence–restraint inversion)**.** *When improved evidence indicates that an intervention has greater potential to produce persistent trajectory change, that evidence can increase reasons for review, reversibility, participation, buffering, or restraint even as it increases confidence about short-run efficacy. More informative evidence can therefore support less direct steering in some near-critical configurations.*

The proposition is deliberately limited. It does not imply that evidence favors inaction or that precaution is always conservative. Evidence of imminent and severe harm can strengthen the case for rapid intervention. The point is that the practical meaning of evidence depends on the regime and on the consequence structure of acting on it. Near criticality, knowledge of leverage is also knowledge of responsibility.

This argument provides a specific extension of the paper's earlier analysis of tangent-space reasoning. Tangent-space information can remain valuable for mapping susceptible directions, identifying uncertainty, locating buffers, and recognizing dangerous couplings. What fails is the automatic conversion of such information into an authorization to exploit the direction of highest leverage. Evidence can inform the boundary of action as much as the choice of action.

17.6 Regime sensitivity and empirical transition

The section yields a provisional limitation of evidence-based policy reasoning: near criticality, evidence-to-action relations can become regime dependent, performative, and dynamically discontinuous. The limitation should be tested rather than assumed. Political systems can display gradual adaptation without critical transition; many policy domains remain sufficiently stable for ordinary evidence transport; and some apparently dramatic episodes can be explained without attractor or tipping language. The framework therefore needs empirical procedures capable of distinguishing ordinary contextual variation from genuine changes in susceptibility and regime structure.

Working Claim 17.7 (Evidence near criticality)**.** Near criticality, evidence should be used to characterize response conditions, uncertainty, transition risk, and the boundaries of warranted action in addition to estimating policy effects. Evidence remains central while its policy role expands beyond selecting the intervention with the largest estimated marginal benefit.

This claim prepares the empirical research program. The next section asks how near-critical political dynamics, noisy recognition, regime-sensitive evidence, response repertoires, reversibility, and transition justice might be studied without presuming the phenomena that the framework is intended to investigate.

18 Empirical Research Program

This section translates the conceptual framework into a research program. Its role is methodological rather than confirmatory. The preceding sections have introduced near criticality, noisy recognition, perturbation–response relations, trajectory closure, response repertoires, embedded governance, reversibility, option space, transition justice, and regime-sensitive evidence. An empirical program should therefore be capable of discovering that some of these constructs are absent, weak, differently organized, or poorly measured in a given political domain. The objective is to identify observable implications, competing explanations, and conditions under which criticality language should be retained, revised, or withdrawn.

The program begins from three forms of discipline. First, a critical transition should not be inferred from historical drama alone. Second, early-warning indicators should be interpreted as conditional evidence rather than as a universal diagnostic code, because variance, autocorrelation, recovery rates, and related measures can produce false alarms or remain silent before some transitions (Scheffer, Bascompte, et al. 2009; Boettiger and Hastings 2012; Dakos et al. 2015). Third, evidence of susceptibility should remain analytically separate from the normative warrant to exploit that susceptibility. Empirical work can strengthen or weaken claims about response structure while leaving questions of justice, legitimacy, and acceptable risk open.

18.1 Construct discipline and empirical posture

The empirical unit is not a generic "crisis." The primary object is a configuration in which response properties appear to be changing: perturbations of comparable scale produce different consequences, recovery slows, coupling patterns reorganize, feedback becomes more self-reinforcing, or formerly accessible trajectories become more difficult to recover. These observations can motivate a near-critical hypothesis. They do not establish one on their own.

Let Z_t denote a latent response regime and let Y_t denote a vector of observables such as event sequences, policy responses, network positions, recovery times, expectation measures, institutional changes, and distributional outcomes. Empirical analysis seeks mappings of the form

$$Y_{1:T} \rightsquigarrow \widehat{Z}_{1:T}, \tag{84}$$

where the arrow denotes fallible inference rather than identification by fiat. The inferred regime should earn its interpretation through converging evidence, model comparison, and the failure of simpler explanations where appropriate.

Research Obligation 18.1 (Construct validity before criticality attribution)**.** A study that attributes near criticality to a political configuration should specify the response property being inferred, the observable indicators used, the relevant baseline, plausible alternative explanations, and the empirical pattern that would count against the criticality interpretation.

This obligation follows the caution developed in Sections 2 and 6. A shift in rhetoric, volatility, or institutional conflict can be politically important while remaining adequately explained by ordinary adaptation, strategic bargaining, exogenous shock, or measurement change. Criticality becomes useful only when it adds discriminating structure to the explanation.

Table 8 summarizes the empirical architecture. The table is intentionally plural. A near-critical hypothesis should ordinarily be stronger when several construct families move in mutually interpretable ways than when one indicator is elevated in isolation.

Table 8: Empirical construct families for near-critical political analysis. The rows identify observable families and inferential cautions rather than a universal measurement protocol.

Construct family	Illustrative observables	Primary empirical caution
Susceptibility	finite response to comparable shocks; response elasticities; recovery time; cross-domain spillover	response depends on direction, scale, baseline, and measurement window
Noise structure	residual variance; temporal correlation; clustered disturbances; cross-unit covariance; event provenance	apparent noise can mix measurement error, unresolved structure, endogenous fluctuation, and external forcing
Coupling	network connectivity; co-movement; interdependence; synchronized response; transmission paths	observed association need not establish causal propagation
Trajectory closure	rising reversal cost; shrinking institutional alternatives; increasing returns; path-dependent commitments	persistence can arise without a critical transition and can differ across actors
Response repertoire	number and diversity of available response modes; switching latency; reversibility; redundancy	nominal options can be inaccessible in practice or correlated under stress
Governance footprint	direct policy effect; expectation shifts; secondary responses; coupling changes; option-space effects	intervention and announcement can jointly alter the field being measured
Transition justice	incidence of benefits and burdens; participation; exit and voice; reversibility; intertemporal effects	normative significance cannot be reduced to one scalar welfare measure

18.2 Response profiles and susceptibility estimation

The first empirical task is to characterize how a political system responds to perturbations of different direction, magnitude, location, and timing. Section 4 introduced response fields as a conceptual alternative to evaluating perturbations by magnitude alone. Empirically, the same idea can be implemented through repeated-response profiles.

Suppose comparable disturbances $\eta^{(k)}$ are observed across time or across matched units. A finite response profile can be written as

$$\mathcal{R}_t = \left\{ \Delta Y_{t+h}^{(k)} \mid \eta_t^{(k)},\, h \in \mathcal{H} \right\}, \tag{85}$$

where $\mathcal{H}$ is a set of observation horizons. The relevant empirical question concerns whether the response profile changes systematically as the configuration evolves. Stronger amplification, longer recovery, wider cross-domain propagation, or greater directional asymmetry can each support a susceptibility hypothesis when alternative mechanisms are examined.

Linear-response ideas and network susceptibility provide useful mathematical comparisons for this task, while the manuscript does not assume that political systems satisfy their physical premises (Kubo 1982; Manik et al. 2017; Bar-Yam and Epstein 2004). Likewise, slowing recovery before some ecological transitions motivates the measurement of recovery dynamics without licensing direct political identification (Nes and Scheffer 2007; Scheffer, Bascompte, et al. 2009).

Hypothesis 18.2 (Response amplification)**.** *For political configurations genuinely approaching a transition characterized by rising susceptibility, comparable perturbations in the relevant response direction should, on average and conditional on the mechanism under study, produce larger or more persistent responses as the transition region is approached.*

The hypothesis is intentionally conditional. A system can approach a transition without monotonic amplification in every observable. Nonstationary forcing, strategic anticipation, institutional buffering, or changes in measurement can all complicate the pattern. Failure to observe amplification across well-chosen response variables therefore weakens one version of the near-critical account while leaving other mechanisms open.

18.3 Noise characterization and perturbation provenance

The second task is to examine what has been classified as noise. Section 3 distinguished measurement error, unresolved structure, stochastic forcing, and potentially generative fluctuation. Empirical work can operationalize this distinction through provenance and dependence analysis.

For each disturbance family, researchers can record at least four properties: source attribution, temporal dependence, cross-unit dependence, and relation to subsequent organization. A disturbance that appears idiosyncratic at one scale may be correlated at another. A residual that appears random under one model may become structured after adding a network, institutional, or temporal variable. The empirical aim is therefore to compare competing decompositions rather than to designate an unexplained component permanently as noise.

Hypothesis 18.3 (Noise reclassification)**.** *In some near-critical political configurations, a subset of observations first classified as residual noise will become more predictable or structurally interpretable when changes in coupling, common exposure, or endogenous feedback are incorporated into the model.*

This hypothesis can fail. If residual structure remains weak across model families and measurement scales, the evidence can support a genuinely stochastic interpretation. Constructive-noise and noise-induced-transition research in other dynamical domains establishes the plausibility of nontrivial noise roles while leaving the political existence and frequency of such roles an empirical question (Astumian and Moss 1998; Forgoston and Moore 2018; Gammaitoni et al. 1998).

18.4 Regime identification and transition inference

The third task concerns regime inference. Political research often observes one historical realization, limited repetitions, and changing measurement systems. The empirical burden is therefore higher than locating a visually dramatic turning point. Candidate transitions should be compared against models of gradual change, structural break without self-reinforcement, externally forced change, strategic coordination, and institutional replacement.

Several empirical signatures can be combined without treating any one as necessary or sufficient: changes in response elasticity, slowing recovery, multimodality or state occupancy changes, increasing reversal cost, feedback that strengthens the new configuration, and asymmetry between forward and return paths. Ecological work on alternative stable states and critical transitions provides a vocabulary for these signatures while also illustrating how demanding their empirical identification can be (Scheffer, Carpenter, et al. 2001; Beisner et al. 2003; Dakos et al. 2015). Political path-dependence research supplies a separate lineage for increasing returns, sequence effects, and rising reversal costs (Pierson 2000; Pierson 1993).

Hypothesis 18.4 (Return asymmetry)**.** *When a political transition involves self-reinforcing trajectory closure rather than a temporary displacement, reversing the initial perturbation or policy change should generally recover the prior configuration less readily than a symmetric-response model would predict.*

Return asymmetry can be studied through interrupted sequences, policy reversal, institutional rollback, matched historical episodes, or simulation calibrated to observed interaction structures. Evidence of rapid symmetric recovery would count against strong hysteresis claims for the case under study.

18.5 Comparative designs and triangulation

No single empirical design is likely to carry the entire research program. Near-critical political claims concern temporal sequence, response sensitivity, network coupling, institutional meaning, reflexive interpretation, and normative incidence. A multi-method architecture is therefore appropriate.

Historical sequence analysis can reconstruct perturbations, institutional responses, feedback loops, and reversal attempts. Comparative case analysis can examine configurations that experienced similar disturbances but different propagation or recovery. Time-series and panel designs can estimate changes in response profiles and recovery dynamics. Network analysis can identify changing transmission pathways and correlated-response vulnerability. Survey and experimental designs can study expectation effects, perceived legitimacy, and behavioral response to warnings or policy signals. Agent-based and other simulation approaches can examine whether proposed mechanisms are capable of producing the qualitative dynamics observed, while remaining clearly separated from empirical validation of the model's assumptions.

Triangulation is especially important because the data-generating process can change during the episode. A single model fitted across the entire sequence may average across regimes that should be distinguished. Conversely, flexible regime-switching models can overfit ordinary variation. The empirical program therefore needs prospective criteria, out-of-sample comparison where feasible, and explicit accounting for model multiplicity.

Research Obligation 18.5 (Competing-model comparison)**.** Empirical studies should compare near-critical explanations with at least one plausible noncritical alternative and should report which observations meaningfully discriminate between them. Model flexibility should not substitute for mechanism-specific evidence.

18.6 Embedded intervention and governance footprints

Section 8 argued that political intervention is itself part of the perturbation field. This claim has direct empirical implications. Evaluation should record the governance footprint beyond the direct policy outcome: expectation changes, anticipatory behavior, secondary responses, network rewiring, distributional shifts, and changes in future option space.

A minimal empirical decomposition can be written as

$$\Delta Y = \Delta Y_{\text{direct}} + \Delta Y_{\text{expectation}} + \Delta Y_{\text{relational}} + \Delta Y_{\text{secondary}} + \varepsilon, \tag{86}$$

where the terms represent analytic components rather than an assumed additive causal model. The purpose is to prevent the direct instrument effect from exhausting the evaluation target.

Hypothesis 18.6 (Intervention performativity)**.** *When actors anticipate a near-critical intervention and can strategically adapt, public signals about the intervention or its evidential basis will sometimes produce measurable response changes before the instrument's direct material effect occurs.*

Policy-feedback research makes this possibility plausible in political systems, where rules and programs can reshape resources, incentives, information, and interpretation (Pierson 1993). The hypothesis remains domain-specific: some interventions are obscure, weakly anticipated, or institutionally insulated and can display little measurable performativity.

18.7 Response repertoires and institutional experiments

The framework also generates empirical questions about institutional architecture. Section 9 proposed that adaptive capacity depends partly on the availability and switchability of heterogeneous response modes. Section 16 added distributed sensing, firebreaks, differentiated latency, review, and reopening

capacity. These propositions can be examined without assuming that decentralization or polycentricity is universally superior.

Relevant measures include the number of practically available response modes, the correlation among response centers, switching time between modes, decision latency under stress, redundancy of sensing channels, burden distribution, and whether review procedures can genuinely reverse or modify an earlier decision. Adaptive-governance and polycentric-governance research provides institutional comparisons for such designs (Folke et al. 2005; Ostrom 2010; Duit and Galaz 2008; May 2022).

Hypothesis 18.7 (Correlated-response vulnerability)**.** *Configurations in which nominally separate decision centers rely on highly correlated information and activate highly synchronized responses will be more vulnerable to system-wide overreaction than configurations with comparable coordination capacity and greater epistemic or response differentiation.*

This hypothesis should be evaluated together with coordination failure. Excess fragmentation can delay protective action or shift burdens onto less powerful actors. The empirical task concerns the relation between response diversity and systemic propagation rather than a generic preference for decentralization.

18.8 Option-space and transition-justice measurement

The normative sections of the paper also create empirical obligations. Option space, agency, reversibility, and transition justice should not remain purely rhetorical labels. Their measurement will often be partial, yet several observable dimensions can be tracked.

For actor or group i, an empirical option profile can include accessible institutional alternatives, exit costs, voice channels, legal or material barriers to reversal, time available for adaptation, and dependence on a narrow set of intermediaries. Changes in these dimensions can be compared before and after a transition or intervention. The resulting profile does not measure justice directly. It identifies consequences that normative analysis must then interpret.

Similarly, transition-justice analysis can document who received benefits, who absorbed risk, whose participation was consequential, which burdens were delayed, and which actors retained meaningful routes for contestation or repair. These observations can be joined to the political-legitimacy, distributive justice, capability, and intergenerational considerations developed in Sections 12–15.

Research Obligation 18.8 (Distribution-aware transition analysis)**.** Empirical claims about successful near-critical response should report the distribution of benefits, burdens, option changes, and reversibility where these consequences are material. Aggregate stabilization alone is insufficient as a complete empirical description of the transition.

18.9 Falsification conditions and negative findings

A useful research program needs outcomes that would weaken its central claims. Near-criticality should therefore remain a defeasible interpretation rather than a retrospective label attached to any major political change.

Several negative findings would matter. Stable response elasticities across the proposed critical interval would weaken susceptibility claims. Rapid and approximately symmetric recovery after reversal would weaken strong hysteresis accounts. Little change in propagation structure would weaken coupling-based transition explanations. Absence of expectation effects in settings where performativity was predicted would weaken the governance-footprint mechanism. Comparable outcomes across systems with substantially different response repertoires would weaken claims about repertoire diversity. Persistent option availability after an alleged trajectory closure would weaken closure claims.

Proposition 18.9 (Empirical defeasibility)**.** *The near-criticality framework gains explanatory value only when its central constructs generate discriminating empirical expectations. Repeated failure of susceptibility, propagation, return-asymmetry, response-repertoire, or trajectory-closure predictions under suitable observations should narrow the framework's domain or motivate conceptual revision.*

This proposition protects the framework from becoming an all-purpose metaphor. Criticality language is most useful when it creates empirical risk for the theory.

18.10 Ethical limits of empirical probing

Some dynamical properties are easiest to identify by perturbing a system. Real political systems create ethical limits on such probing. Deliberately increasing social instability, manipulating vulnerable populations, or exploiting a suspected transition window for research purposes can impose harms far beyond the informational benefit of the study. The research program therefore favors observational perturbations, naturally occurring variation, low-risk laboratory or survey settings, archival reconstruction, simulation, and institutionally approved reversible pilots where these designs can answer the question.

Research Obligation 18.10 (Non-escalatory empirical design)**.** Empirical investigation of suspected near-critical political configurations should avoid research interventions whose informational purpose depends on creating a material risk of irreversible political transition, concentrated harm, or loss of agency for affected populations.

This obligation mirrors the paper's substantive argument. Greater leverage can increase the responsibility attached to its use. The same principle applies to research design.

18.11 Sequenced research agenda

The empirical program can proceed in stages. A first stage develops case-specific response maps and noise characterizations. A second stage tests whether susceptibility, recovery, coupling, and reversal properties change over time. A third stage compares near-critical and ordinary-regime explanations across multiple cases or simulated mechanisms. A fourth stage examines institutional response repertoires, governance footprints, option-space effects, and transition justice. A fifth stage evaluates whether the same constructs travel across domains or whether several distinct families of political criticality are needed.

The sequence leaves room for substantial revision. Some domains may support strong dynamical identification. Others may support only a weaker vocabulary of heightened susceptibility and path dependence. Still others may be better explained without criticality language. Such variation would be a result rather than a failure of the program.

Working Claim 18.11 (Empirical orientation)**.** The empirical value of near-criticality analysis lies in identifying changes in response structure, propagation, recovery, reversibility, and reachable trajectories while preserving uncertainty about mechanism, normative meaning, and cross-domain generality. The research program therefore treats criticality as a hypothesis about political dynamics whose domain must be earned case by case.

This orientation prepares the remaining open problems. The next section turns to unresolved questions concerning scale, cross-domain transport, endogenous normativity, observer participation, historical contingency, and the limits of the criticality vocabulary itself.

19 Open Problems in Near-Critical Political Response

This section identifies unresolved problems that remain after the conceptual, normative, institutional, and empirical arguments developed above. Its role is to mark the present boundary of the framework rather than to append a generic list of future research topics. Near-critical political response combines several forms of uncertainty that cannot yet be resolved by one vocabulary: uncertainty about scale and system boundaries, uncertainty about whether criticality mechanisms travel across political domains, uncertainty about what counts as noise, reflexive changes produced by observation, normative pluralism, historical contingency, and strategic uses of criticality language itself. Each problem below can alter the interpretation of susceptibility, response, trajectory closure, or legitimate intervention.

The open problems are connected. A scale choice changes the relevant noise process; a noise model changes estimates of susceptibility; susceptibility claims can become public information; public information changes strategic behavior; strategic behavior changes the response field; and the resulting transition can alter the institutions through which justice and legitimacy are subsequently evaluated. The research task therefore concerns the mutual relations among these problems in addition to their separate treatment.

Table 9 summarizes the section. The common concern is epistemic and political restraint: each unresolved problem identifies a place where analytical confidence can exceed what the current evidence or normative argument supports.

19.1 Scale, Boundary, and Unit Selection

Near criticality is always attributed relative to a system description. A configuration that appears highly susceptible at one level can remain buffered at another. A local institution can cross a transition while a broader network absorbs the change; a global indicator can appear stable while particular communities experience severe and persistent closure. Temporal scale creates a similar problem. A disturbance that looks transient over weeks can restructure commitments over years, while a dramatic short-run shift can later dissipate.

Let $S_t^{(\ell)}$ denote a system representation at scale ℓ. Susceptibility is then better written as

$$\chi_t^{(\ell,h,d)} = \frac{\partial R_{t+h}^{(\ell)}}{\partial \eta_t^{(d)}}, \tag{87}$$

where h is the response horizon and d denotes a perturbation direction. Equation (87) makes explicit that a claim about rising susceptibility is incomplete until the relevant level, horizon, and direction are specified.

Political systems also contain overlapping boundaries. Trade, security, finance, law, infrastructure, public opinion, ecological dependence, and information networks can connect different sets of actors. A transition can therefore propagate across boundaries that were analytically separate at the start of the study. This creates a moving-boundary problem: the system relevant to explanation can expand as coupling changes.

Research Obligation 19.1 (Scale declaration)**.** Near-criticality claims should state the unit of analysis, observation horizon, perturbation direction, and principal coupling boundary. Cross-scale effects should be treated as empirical questions when they materially affect the interpretation of transition or recovery.

The obligation does not privilege one scale. Its purpose is to prevent a politically convenient scale from silently determining what counts as stability, harm, or successful response.

19.2 Criticality Families and Cross-Domain Transport

The manuscript has used criticality as a family resemblance among conditions of heightened susceptibility, nonlinear response, changing recovery, feedback, and possible regime transition. A remaining prob-

Table 9: Principal open problems for near-critical political response. The matrix records unresolved analytical tasks rather than a hierarchy of research importance.

Problem domain	**Unresolved analytical task**	**Risk of premature closure**
Scale and boundary	identify relevant units, levels, time horizons, and coupling boundaries	local stability or instability can be mistaken for a system-level property
Criticality families	distinguish mechanisms that share sensitivity language while differing dynamically	a single metaphor can conceal several transition processes
Noise ontology	separate stochasticity, measurement error, omitted structure, and emergent organization	residuals can be suppressed or interpreted before their role is understood
Reflexivity	model how observation, warning, and intervention alter expectations and behavior	evidence can become a perturbation to the system it describes
Normative evolution	evaluate transitions when values, identities, and institutions also change	present criteria can be imposed as though historically invariant
Power and representation	identify who defines thresholds, risks, acceptable losses, and response options	technical vocabulary can legitimate asymmetric trajectory control
Historical contingency	determine how prior sequences shape present susceptibility and reachable futures	comparable states can be treated as dynamically equivalent despite different histories
Conceptual translation	use comparative resources such as *wuwei* and *shi* with historical discipline	analogy can flatten philosophical traditions into governance techniques
Domain limitation	specify when criticality language should be narrowed or abandoned	the framework can become unfalsifiable through retrospective relabeling

lem concerns the internal heterogeneity of that family. Fold-like transitions, noise-induced switching, network cascades, strategic coordination shifts, institutional path dependence, and expectation-driven transitions can share some observable features while requiring different causal explanations. Early-warning indicators developed for one mechanism can therefore travel poorly to another (Scheffer, Bascompte, et al. 2009; Boettiger and Hastings 2012; Dakos et al. 2015).

The political problem is larger because reflexive agency can change both parameters and rules. Actors can learn from earlier crises, redesign institutions, conceal intentions, coordinate strategically, or deliberately alter network structure. A political transition may therefore combine several mechanisms whose relative importance changes during the episode.

Working Claim 19.2 (Plural criticality families)**.** A mature theory of near-critical political response may require several mechanism-specific families rather than a single general model. Shared concepts such as susceptibility, amplification, return asymmetry, and option contraction can organize comparison while leaving mechanism-specific predictions distinct.

This claim leaves open how many families are useful. The empirical program in Section 18 can help determine whether a common vocabulary provides explanatory compression or whether it masks differences that matter for response and ethics.

19.3 Noise Ontology and Model Dependence

Section 3 treated random noise as a dialectical object. The open problem is deeper than classification accuracy. The category of noise partly depends on the model through which a system is observed. A residual process can represent measurement error, genuine stochastic forcing, strategic behavior omitted from the model, cross-scale influence, or an emerging pattern whose structure has not yet been recognized. Model revision can consequently move observations between "signal" and "noise" without any change in the underlying political world.

A useful formal reminder is

$$Y_t = M_k(X_t) + \varepsilon_t^{(k)}, \tag{88}$$

where the residual $\varepsilon_t^{(k)}$ is indexed by model M_k. Different models generate different residuals. The residual therefore should not be identified immediately with an ontologically independent random force.

At the same time, genuine stochasticity remains possible. Constructive-noise, stochastic-resonance, and noise-induced-transition literatures demonstrate that random forcing can sometimes alter response and transition behavior in other dynamical systems (Gammaitoni et al. 1998; Astumian and Moss 1998; Forgoston and Moore 2018). The political existence, frequency, and normative significance of analogous roles remain empirical questions.

Remark 19.3 (Epistemic status of noise)**.** The phrase "epistem the noise" names a research posture rather than a promise that every fluctuation conceals meaningful structure. The posture requires examining provenance, dependence, scale, and response before assigning a stable political role to the residual.

A further open question concerns institutional incentives. Organizations often benefit from classifying inconvenient variation as noise or, conversely, from presenting ordinary variation as evidence of exceptional instability. Noise classification can therefore become part of political contestation.

19.4 Reflexivity and Observer Participation

Near-critical political knowledge can become performative. Publication of a warning, forecast, stress assessment, intelligence judgment, or policy model can change expectations before any formal intervention occurs. Actors can hedge, exit, coordinate, mobilize, imitate, or strategically counter the expected response. The observer is therefore potentially part of the perturbation field.

Let O_t denote an observation or public assessment and E_t denote the expectation state of relevant actors. A reflexive sequence can be represented as

$$O_t \mapsto E_{t+1} \mapsto A_{t+1} \mapsto G_{t+1}, \tag{89}$$

where A_{t+1} denotes actions and G_{t+1} the resulting coupling structure. The observation can therefore change the object whose future was being estimated.

This problem complicates open science, public warning, and democratic accountability. Withholding information can concentrate power and reduce public agency. Immediate disclosure can also alter a fragile field. A general rule of secrecy or disclosure would reproduce the same premature closure criticized elsewhere in the paper.

Research Obligation 19.4 (Reflexive disclosure analysis)**.** Where criticality assessments are plausibly performative, researchers and institutions should examine the likely effects of disclosure, delay, aggregation, uncertainty communication, and selective release together with the rights and interests of affected publics.

The open problem concerns institutional forms capable of supporting this analysis without converting epistemic caution into unaccountable secrecy.

19.5 Normative Evolution and Transition Standing

The normative sections of the paper have treated justice, legitimacy, agency, reversibility, and option space as dimensions through which transitions can be evaluated. A difficult unresolved problem appears when the transition changes the evaluative field itself. Political transformations can alter institutions, identities, preferences, legal categories, social expectations, and the groups recognized as legitimate participants. The standards used before a transition and the standards endorsed after it can therefore diverge.

Let N_t denote a normative-evaluative configuration. A transition then has a joint structure

$$(X_t, N_t) \longmapsto (X_{t+1}, N_{t+1}), \tag{90}$$

where political state and evaluative framework can co-evolve. The notation does not imply that justice is reducible to social preference. It records the historical fact that the institutions through which claims are articulated and recognized can change during the process being judged.

Several tensions follow. Present actors can impose current norms on futures whose participants may understand their interests differently. Future approval can also fail to legitimate a transition that achieved compliance through coercion, exclusion, or destruction of alternative identities. Intergenerational justice and capability approaches already give reasons to examine future opportunities and unequal conversion of resources into real possibilities (Meyer 2021; Robeyns and Byskov 2025); near-critical analysis adds the possibility that the evaluative institutions themselves are part of the transition.

Working Claim 19.5 (Normative co-evolution)**.** Evaluation of near-critical transitions should distinguish changes in political state from changes in the normative and institutional conditions through which that state is later evaluated. Ex post stabilization or endorsement does not by itself settle the justice of the transition path.

Transition standing remains open within this problem. Those exposed to risks, those whose options are closed, future persons, external populations affected by spillovers, and institutions charged with collective responsibility can each have different claims to participation. The framework currently identifies these dimensions without providing one final rule for resolving conflicts among them.

19.6 Aesthetic Salience and Normative Capture

Section 14 argued that aesthetic judgment can register an emerging form before its causal and normative meaning is fully articulated. The open problem concerns the governance of that sensitivity itself. Political actors can be attracted to coherence, symmetry, simplicity, unity, momentum, or narrative elegance. These qualities can aid recognition while also encouraging premature commitment to one emerging order.

The problem becomes acute near criticality because an aesthetically salient configuration can attract attention precisely when the field is highly susceptible. Recognition can then become advocacy, and advocacy can become a perturbation that helps constitute the order initially perceived only as a possibility. The distinction among aesthetic salience, empirical plausibility, and normative warrant therefore requires institutional support in addition to conceptual clarity.

Research Obligation 19.6 (Aesthetic pluralization)**.** Where emerging political forms are evaluated under substantial uncertainty, analysis should preserve space for multiple interpretations of coherence, fragmentation, novelty, and disorder, especially when aesthetic descriptions could influence trajectory selection.

The remaining theoretical task is to understand how aesthetic judgment can support inquiry and imagination while limiting aesthetic capture. This problem connects political aesthetics to the paper's broader concern with delayed normative closure (Ginsborg 2022; Tömmel and Passerin d'Entreves 2025; Rancière 2013).

19.7 Historical Contingency and Path Reconstruction

Near-critical response is history dependent. Two systems with similar current observables can possess different latent commitments, memories, institutional routines, trust relations, and reversal costs because they arrived through different sequences. Path-dependence research gives a vocabulary for increasing returns and self-reinforcing institutional development (Pierson 2000; Pierson 1993). The open problem concerns how much historical depth is required to distinguish genuine dynamical similarity from superficial state similarity.

A state-only representation can write

$$X_t^{(a)} \approx X_t^{(b)}, \tag{91}$$

while a trajectory-sensitive representation records

$$H_{0:t}^{(a)} \neq H_{0:t}^{(b)}, \qquad \mathcal{R}_t^{(a)} \neq \mathcal{R}_t^{(b)}, \tag{92}$$

where $H_{0:t}$ denotes relevant historical development and $\mathcal{R}_t$ the current response profile. Similar observable states can therefore coexist with different response fields.

This creates a reconstruction problem. Historical records are selective, institutionally produced, and often uneven across groups. The trajectories most important for explaining present susceptibility can be precisely those least visible in official archives. Empirical work therefore needs methods that combine temporal depth with source criticism and distributional awareness.

19.8 Institutional Power and Strategic Criticality Claims

Criticality language can itself become a political resource. An actor who successfully declares a system to be near collapse can seek exceptional powers, accelerated procedures, secrecy, emergency coordination, suppression of variation, or tolerance for concentrated losses. Another actor can strategically deny susceptibility in order to delay costly adaptation. The framework must therefore analyze claims about criticality together with the institutions and interests through which those claims acquire authority.

This problem extends the legitimacy discussion in Section 12. Expertise about susceptibility creates epistemic asymmetry, while political action based on that expertise can redistribute rights, burdens, and future options. The relevant issue is not a generic opposition between expertise and democracy. It concerns the procedures through which uncertain high-leverage knowledge is challenged, communicated, and converted into action.

Proposition 19.7 (Criticality-claim accountability)**.** *Institutions that invoke near-criticality to justify exceptional or trajectory-shaping action should expose the relevant indicators, uncertainty, assumptions, expected distributional consequences, review conditions, and reopening mechanisms to forms of scrutiny proportionate to the stakes and security constraints of the case.*

The proposition remains intentionally procedural. It cannot by itself settle which emergency powers, confidentiality rules, or response thresholds are legitimate. Those questions require domain-specific law, ethics, and political judgment.

19.9 Comparative Conceptual Translation

The paper has used *wuwei*, water imagery, and *shi* as comparative resources for thinking about non-coercive response and configuration-dependent propensity. Their continued use creates an interpretive obligation. These terms come from heterogeneous Chinese philosophical and strategic traditions, and their meanings vary across texts, periods, and contemporary scholarship (Chan 2025; Slingerland 2003; Jullien 1995; Raphals 2016; Yue et al. 2023).

The open problem concerns productive translation without conceptual capture. If *wuwei* is reduced to "minimal intervention," much of its ethical and philosophical context disappears. If *shi* is reduced to a scalar measure of leverage, the term adds little beyond existing dynamical vocabulary. Their value in this manuscript lies partly in the resistance they offer to a simple controller–object grammar. Future development should therefore compare the concepts at the level of problems and relations rather than claiming direct equivalence.

Research Obligation 19.8 (Comparative concept discipline)**.** Comparative philosophical concepts used in near-criticality analysis should be accompanied by explicit limits on translation, relevant textual or scholarly support, and a statement of which analytical problem the comparison clarifies.

This obligation also applies in the opposite direction. Complexity vocabulary should not be projected back onto historical philosophical traditions as though they had anticipated contemporary dynamical systems theory.

19.10 Domain Limits and Theory Exit Conditions

The final open problem concerns the limits of the framework itself. Criticality is attractive because it offers a language for disproportionate response, transition, feedback, and irreversibility. Those same features make it easy to extend the vocabulary beyond its explanatory domain. A theory that can redescribe every political crisis, reform, conflict, or institutional change as near-critical loses discriminating value.

Section 18 proposed empirical falsification conditions. The conceptual framework also needs exit conditions. Criticality language should recede when ordinary bargaining, linear adaptation, exogenous forcing, gradual institutional change, or other established explanations account for the observed response structure with fewer unsupported assumptions. Likewise, concepts such as attractor, hysteresis, and trajectory closure should be narrowed when persistence, return asymmetry, or constrained reachability cannot be demonstrated at the relevant scale.

Proposition 19.9 (Theory exit condition)**.** *Near-criticality analysis should be suspended or narrowed when the framework fails to add discriminating explanatory structure beyond plausible alternative accounts, or when its central response and transition constructs cannot be connected to observable consequences at the scale under study.*

The proposition gives conceptual restraint a positive role. Withdrawal of the criticality vocabulary can represent successful inquiry when the evidence supports a simpler account.

19.11 Research Frontier and Conceptual Provisionality

The open problems collectively place the framework in a provisional position. The manuscript has developed a vocabulary for susceptibility, noise, perturbation–response relations, trajectory closure, response repertoires, non-coercive participation, propensity, transition justice, reversibility, option space, and regime-sensitive evidence. The remaining questions concern how these concepts interact across scales, histories, institutions, and normative orders.

A particularly important frontier is the relation between response and historical creation. Near-critical configurations can make several futures newly reachable while leaving their meaning incompletely manifested. Political actors can participate in shaping those futures, yet the standards by which the result is later judged can themselves evolve. The framework therefore cannot promise a timeless decision rule for critical moments. Its present contribution is narrower: to make visible the dynamical, epistemic, normative, and institutional burdens that arise when small responses can participate in large and persistent transitions.

Working Claim 19.10 (Provisional research horizon)**.** The theory remains most useful when it preserves the distinction between recognizing heightened susceptibility and claiming authority over the future. Its research horizon concerns the conditions under which political actors can observe, respond, protect, experiment, deliberate, and revise while acknowledging that near-critical transitions can exceed both prediction and intention.

This horizon prepares the conclusion. The final section consolidates the paper's principal distinctions and clarifies the limited sense in which "governance" remains an appropriate term for responsive participation near criticality.

20 Conclusion

This paper has examined a problem that becomes unusually difficult when a political system approaches a region of heightened susceptibility: how should political actors respond when comparatively small perturbations can participate in large and persistent changes, while the relevant transition mechanisms, future consequences, and normative standing of available trajectories remain only partially known? The paper began from the language of criticality, but its argument has progressively displaced a simple control-theoretic interpretation. Near criticality, the central problem is not merely how to identify a sensitive direction and exploit it efficiently. It is how to understand perturbations, response fields, trajectory closure, and political responsibility when the act of intervention can itself become part of the dynamics under examination.

The resulting framework is deliberately conditional. It does not assume that international or political systems possess one universal critical point, that every abrupt change is a phase transition, or that concepts such as attractor, hysteresis, and susceptibility can be transferred from physical systems without empirical reconstruction. Section 2 therefore treated near criticality as a domain claim that must be specified by scale, horizon, direction, coupling structure, and observable response. Sections 18 and 19 subsequently placed falsification conditions and exit conditions on the theory itself. The framework is useful only where it adds discriminating structure to the analysis of political response.

20.1 From perturbation control to response analysis

The first major shift concerns the status of noise. Section 3 argued that random or residual variation should not be treated automatically as an error term awaiting elimination. What appears as noise can include measurement error, unresolved structure, stochastic forcing, endogenous fluctuation, or the early manifestation of a new organization. Near criticality, these possibilities become politically significant because a fluctuation of small initial magnitude can interact with susceptibility and coupling in ways that exceed its apparent scale.

The paper therefore moved from the isolated perturbation η_t to the perturbation–response relation developed in Section 4. The relevant political object is not only what disturbance occurs, but what the receiving configuration does with it. A perturbation can be absorbed, redirected, delayed, propagated, amplified, or incorporated into a larger transformation. Its significance depends on the response field, historical state, network structure, and the scales across which the response travels. This is the analytical basis for the paper's recurring formulation of "episteming the noise": the first task is to investigate the role of a fluctuation before deciding that its appropriate treatment is suppression.

This perspective changes the meaning of adaptive capacity. Section 9 argued that capacity near criticality is partly a property of the diversity, differentiation, and revisability of available responses. A system that can only intensify control may possess considerable coercive power while remaining poorly equipped for an uncertain response field. Observation, buffering, decoupling, delay, reversible experimentation, redirection, selective amplification, restraint, and stronger protective action can each have a place within a response repertoire. Their appropriateness remains configuration dependent.

Working Claim 20.1 (Response-centered orientation)**.** Near-critical political analysis is better organized around the relation between perturbation and response than around a prior assumption that perturbations should be eliminated or exploited. This orientation keeps open the possibility that fluctuations can reveal susceptibility, support adaptation, or contribute to emerging organization while retaining the possibility that some perturbations require rapid containment.

The claim does not romanticize noise. Catastrophic cascades, irreversible harm, and concentrated exposure remain possible. The shift concerns analytical order: classification and response should follow investigation of the relation rather than precede it through a general presumption that variance is pathological.

20.2 Limits of steering under susceptibility

The second major shift concerns the meaning of leverage. In ordinary policy reasoning, evidence that a small intervention can produce a large desired effect may increase the attractiveness of intervention. Near criticality, the same finding has a second interpretation. High leverage also means that error, misclassification, strategic reaction, and normative disagreement can acquire larger consequences. Sections 7 and 17 therefore separated local epistemic validity, dynamical reversibility, and normative warrant.

This separation is central to the paper. Tangent-space information can remain useful for identifying susceptible directions even when acting along those directions becomes dangerous or politically unjustified. Likewise, evidence can remain statistically or causally informative while its transport into a near-critical regime becomes uncertain. A policy can also alter the regime in which the evidence was generated, change actors' expectations, and produce secondary behavior that was absent from the original evidence base. The relation between evidence and action is therefore itself regime sensitive.

Section 8 extended this argument by placing governance inside the perturbation field. An intervention is an intentional act for its initiator and simultaneously a perturbation for other actors. Their responses can generate further disturbances, expectations, alignments, and feedbacks. The resulting governance footprint can exceed the intended policy channel. This recursive structure weakens the image of an external controller standing above a political system and strengthens the need for reflexive accounts of intervention.

The paper's use of *wuwei* and *shi* in Sections 10 and 11 was intended to clarify this limit rather than to provide a ready-made doctrine of government. *Wuwei* supplied a comparative vocabulary for non-coercive and non-forcing response, while *shi* helped distinguish configuration-dependent propensity from force possessed by an isolated actor. Their analytical contribution lies in challenging a simple controller–object grammar. A political actor can respond to a field of tendencies without claiming to determine the field's final form.

Proposition 20.2 (Leverage–responsibility separation)**.** *Greater knowledge of near-critical leverage can increase practical efficacy without increasing political entitlement to use that leverage. When small acts can participate in persistent trajectory change, improved knowledge of susceptibility can enlarge responsibilities of justification, review, participation, reversibility, and restraint.*

This proposition captures one of the paper's principal departures from a narrow control interpretation. Near criticality, technical capability and normative authority can diverge.

20.3 Transition ethics and open futures

The third shift concerns what counts as the object of normative evaluation. Sections 12 and 13 distinguished the justice of an outcome from the justice of the path through which that outcome is produced. A transition can generate aggregate benefits while distributing losses, uncertainty, voice, and irreversible exposure unequally. Conversely, restraint can preserve institutional flexibility while allowing harms to accumulate for actors with fewer buffers. The ethical problem therefore includes intervention, delay, and non-intervention.

Near-criticality sharpens these questions because a transition can alter more than present welfare. Section 15 introduced option space, trajectory agency, and political revisability to capture changes in what actors remain able to do later. A response can improve one present indicator while narrowing the set of futures that particular actors can meaningfully reach. The distribution of future revisability can therefore become part of transition justice.

This does not establish openness as an unconditional good. Some options represent severe and continuing harms; some boundaries require urgent protection; and some transitions can expand agency by closing oppressive arrangements. The more limited principle developed here is that option contraction

acquires special normative importance when it is persistent, unequally distributed, weakly authorized, or difficult to repair. This is why reversibility has both dynamical and ethical significance.

Aesthetic judgment, considered in Section 14, adds another layer without resolving the normative problem. Political actors can sometimes perceive an emerging coherence or possible form before its consequences are fully manifested. Such perception can guide attention and inquiry. It cannot establish justice. A form can appear elegant, coherent, or historically promising while imposing severe burdens on actors who experience it differently. The paper therefore kept aesthetic salience, dynamical plausibility, and normative warrant analytically distinct.

Research Obligation 20.3 (Transition responsibility)**.** Intentional responses near criticality should be assessed in relation to the persistence of the transition they may support, the distribution of exposure and option loss, the standing of affected actors, the availability of review or repair, and the uncertainty under which the trajectory is being selected.

Transition responsibility is not a rule for choosing one attractor. It is a constraint on how seriously political actors should treat their participation in trajectory closure.

20.4 Institutional and empirical implications

The framework also produces a different institutional emphasis. Section 16 argued for architectures capable of sensing and responding without converting every warning into synchronized action. Distributed observation, epistemic independence, institutional firebreaks, differentiated response speeds, reversible measures, plural review, and memory can help preserve response capacity under uncertainty. The aim is neither maximal decentralization nor maximal central command. It is to reduce correlated failure while retaining the ability to coordinate when boundary protection becomes urgent.

The empirical program in Section 18 makes these claims testable in principle. Susceptibility, response profiles, coupling, return asymmetry, governance footprints, option contraction, and response repertoires should be connected to observable consequences rather than inferred from dramatic historical outcomes alone. Evidence against criticality should narrow the framework's domain. Evidence for gradual adaptation, ordinary bargaining, or simpler causal mechanisms should remain capable of displacing the near-critical interpretation.

This empirical openness matters because political uses of criticality language can themselves become strategic. Claims that a system is near a tipping point can be used to justify urgency, exceptional action, concentration of authority, or deliberate manipulation of sensitive directions. A theory of near-critical response therefore requires scrutiny of who identifies the threshold, who defines unacceptable trajectories, who bears transition risk, and whose interpretation becomes institutionally authoritative.

20.5 The remaining meaning of governance

The paper retains the word "governance" only in a qualified sense. If governance is defined as selecting a desired state and applying sufficient control to move a system toward it, the concept fits poorly with the framework developed here. Near criticality combines incomplete knowledge, reflexive agents, potentially generative noise, regime-dependent evidence, normative pluralism, and the possibility that intervention itself triggers trajectory closure. Under these conditions, governance cannot plausibly mean sovereign command over the transition landscape.

A more limited meaning remains available. Governance can refer to the situated practice of observing, protecting, buffering, deliberating, experimenting, redirecting, coordinating, restraining, and revising within a field whose final form is not fully known in advance. The word then names participation in response dynamics rather than mastery over them. This interpretation also leaves room for moments in which strong action is justified. Non-coercive response does not imply passivity, and restraint does not

possess automatic moral priority. The relevant question concerns what form of participation is warranted by the configuration, uncertainty, irreversibility, affected agency, and distribution of risk.

Working Claim 20.4 (Near-critical political response)**.** Near-critical political response concerns the responsible participation of political actors in highly susceptible and partially knowable response fields. Its central tasks are to characterize noise and susceptibility, maintain a plural repertoire of responses, protect against catastrophic closure, preserve revisability where possible, examine the justice of transition paths, and recognize that the capacity to steer does not by itself establish the right to steer.

The claim leaves an important question open. History can present political actors with transition possibilities whose meaning cannot be fully evaluated from the standpoint that precedes them. The standards of judgment, identities, and institutions through which a later order is interpreted can themselves evolve during the transition. A theory of near-critical response can therefore clarify burdens of knowledge, action, and justification without supplying a final algorithm for historical becoming.

The paper's concluding orientation is consequently one of disciplined responsiveness. Near criticality can make small acts consequential, but this fact does not convert politics into an optimization problem. It makes the relation among knowledge, intervention, justice, and future possibility more serious. Where leverage is high and reversibility is low, political wisdom may consist partly in knowing how to respond without prematurely converting partial knowledge into irreversible trajectory control. The task is to remain capable of protection and action while preserving, where conditions permit, the space in which collective futures can continue to be observed, contested, revised, and generated.

References

Ashby, W. Ross (1956). *An Introduction to Cybernetics*. London: Chapman & Hall. URL: https://ashby.info/Ashby-Introduction-to-Cybernetics.pdf.

Astumian, R. Dean and Frank Moss (1998). "Overview: The Constructive Role of Noise in Fluctuation Driven Transport and Stochastic Resonance". In: *Chaos* 8.3, pp. 533–538. DOI: 10.1063/1.166334.

Bar-Yam, Yaneer and Irving R. Epstein (2004). "Response of Complex Networks to Stimuli". In: *Proceedings of the National Academy of Sciences* 101.13, pp. 4341–4345. DOI: 10.1073/pnas.0400673101.

Beisner, Beatrix E., Daniel T. Haydon, and Kim Cuddington (2003). "Alternative Stable States in Ecology". In: *Frontiers in Ecology and the Environment* 1.7, pp. 376–382. DOI: 10.1890/1540-9295(2003)001[0376:ASSIE]2.0.CO;2.

Boettiger, Carl and Alan Hastings (2012). "Quantifying Limits to Detection of Early Warning for Critical Transitions". In: *Journal of the Royal Society Interface* 9.75, pp. 2527–2539. DOI: 10.1098/rsif.2012.0125.

Cairney, Paul (2016a). *The Politics of Evidence-Based Policy Making*. London: Palgrave Macmillan. DOI: 10.1057/978-1-137-51781-4. URL: https://doi.org/10.1057/978-1-137-51781-4.

— (2016b). "To Bridge the Divide between Evidence and Policy: Reduce Ambiguity as Much as Uncertainty". In: *Public Administration Review* 76.3, pp. 399–402. DOI: 10.1111/puar.12555. URL: https://doi.org/10.1111/puar.12555.

Cartwright, Nancy and Jeremy Hardie (2012). *Evidence-Based Policy: A Practical Guide to Doing It Better*. Oxford: Oxford University Press. ISBN: 9780199841608. DOI: 10.1093/acprof:osobl/9780199841608.001.0001. URL: https://doi.org/10.1093/acprof:osobl/9780199841608.001.0001.

Chan, Alan (2025). "Laozi". In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta and Uri Nodelman. Winter 2025. Metaphysics Research Lab, Stanford University. URL: https://plato.stanford.edu/archives/win2025/entries/laozi/.

Conant, Roger C. and W. Ross Ashby (1970). "Every Good Regulator of a System Must Be a Model of That System". In: *International Journal of Systems Science* 1.2, pp. 89–97. DOI: 10.1080/00207727008920220.

Dakos, Vasilis, Stephen R. Carpenter, Egbert H. van Nes, and Marten Scheffer (2015). "Resilience Indicators: Prospects and Limitations for Early Warnings of Regime Shifts". In: *Philosophical Transactions of the Royal Society B: Biological Sciences* 370.1659, p. 20130263. DOI: 10.1098/rstb.2013.0263.

Duit, Andreas and Victor Galaz (2008). "Governance and Complexity—Emerging Issues for Governance Theory". In: *Governance* 21.3, pp. 311–335. DOI: 10.1111/j.1468-0491.2008.00402.x.

Folke, Carl, Thomas Hahn, Per Olsson, and Jon Norberg (2005). "Adaptive Governance of Social-Ecological Systems". In: *Annual Review of Environment and Resources* 30, pp. 441–473. DOI: 10.1146/annurev.energy.30.050504.144511.

Forgoston, Eric and Richard O. Moore (2018). "A Primer on Noise-Induced Transitions in Applied Dynamical Systems". In: *SIAM Review* 60.4, pp. 969–1009. DOI: 10.1137/17M1142028.

Gammaitoni, Luca, Peter Hänggi, Peter Jung, and Fabio Marchesoni (1998). "Stochastic Resonance". In: *Reviews of Modern Physics* 70.1, pp. 223–287. DOI: 10.1103/RevModPhys.70.223.

Ginsborg, Hannah (2022). "Kant's Aesthetics and Teleology". In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta and Uri Nodelman. Fall 2022. Metaphysics Research Lab, Stanford University. URL: https://plato.stanford.edu/archives/fall2022/entries/kant-aesthetics/.

Haasnoot, Marjolijn, Jan H. Kwakkel, Warren E. Walker, and Judith ter Maat (2013). "Dynamic Adaptive Policy Pathways: A Method for Crafting Robust Decisions for a Deeply Uncertain World". In: *Global Environmental Change* 23.2, pp. 485–498. DOI: 10.1016/j.gloenvcha.2012.12.006.

Hansson, Sven Ove (2022). "Risk". In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta and Uri Nodelman. Winter 2022. Metaphysics Research Lab, Stanford University. URL: https://plato.stanford.edu/entries/risk/.

Head, Brian W. (2008). "Three Lenses of Evidence-Based Policy". In: *Australian Journal of Public Administration* 67.1, pp. 1–11. DOI: 10.1111/j.1467-8500.2007.00564.x. URL: https://doi.org/10.1111/j.1467-8500.2007.00564.x.

— (2010). "Reconsidering Evidence-Based Policy: Key Issues and Challenges". In: *Policy and Society* 29.2, pp. 77–94. DOI: 10.1016/j.polsoc.2010.03.001. URL: https://doi.org/10.1016/j.polsoc.2010.03.001.

Holling, C. S. (1973). "Resilience and Stability of Ecological Systems". In: *Annual Review of Ecology and Systematics* 4, pp. 1–23. DOI: 10.1146/annurev.es.04.110173.000245.

Jullien, François (1995). *The Propensity of Things: Toward a History of Efficacy in China*. Trans. by Janet Lloyd. New York: Zone Books. ISBN: 9780942299946.

Kharrazi, Ali, Yadong Yu, Arun Jacob, Nemi Vora, and Brian D. Fath (2020). "Redundancy, Diversity, and Modularity in Network Resilience: Applications for International Trade and Implications for Public Policy". In: *Current Research in Environmental Sustainability* 2, p. 100006. DOI: 10.1016/j.crsust.2020.06.001. URL: https://doi.org/10.1016/j.crsust.2020.06.001.

Kubo, Ryogo (1982). "Fluctuation, Response, and Relaxation: The Linear Response Theory Revisited". In: *International Journal of Quantum Chemistry* 22.S16, pp. 25–32. DOI: 10.1002/qua.560220806.

Lamont, Julian and Christi Favor (2017). "Distributive Justice". In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta. Winter 2017. Metaphysics Research Lab, Stanford University. URL: https://plato.stanford.edu/entries/justice-distributive/.

Law, John and Wen-yuan Lin (2018). "Tidescapes: Notes on a Shi-Inflected Social Science". In: *Journal of World Philosophies* 3.1, pp. 1–16. URL: https://scholarworks.iu.edu/iupjournals/index.php/jwp/article/view/1612.

Leddy, Tom and Kalle Puolakka (2025). "Dewey's Aesthetics". In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta and Uri Nodelman. Winter 2025. Metaphysics Research Lab, Stanford University. URL: https://plato.stanford.edu/archives/win2025/entries/dewey-aesthetics/.

Lovett, Frank (2026). "Republicanism". In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta and Uri Nodelman. Fall 2026. Metaphysics Research Lab, Stanford University. URL: https://plato.stanford.edu/archives/fall2026/entries/republicanism/.

Manik, Debsankha, Martin Rohden, Henrik Ronellenfitsch, Xiaozhu Zhang, Sarah Hallerberg, Dirk Witthaut, and Marc Timme (2017). "Network Susceptibilities: Theory and Applications". In: *Physical Review E* 95.1, p. 012319. DOI: 10.1103/PhysRevE.95.012319.

May, Candace K. (2022). "Complex Adaptive Governance Systems: A Framework to Understand Institutions, Organizations, and People in Socio-Ecological Systems". In: *Socio-Ecological Practice Research* 4, pp. 39–54. DOI: 10.1007/s42532-021-00101-7.

Meyer, Lukas H. (2021). "Intergenerational Justice". In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta. Fall 2024 archive, substantive revision 2021. Metaphysics Research Lab, Stanford University. URL: https://plato.stanford.edu/archives/fall2024/entries/justice-intergenerational/.

Nes, Egbert H. van and Marten Scheffer (2007). "Slow Recovery from Perturbations as a Generic Indicator of a Nearby Catastrophic Shift". In: *The American Naturalist* 169.6, pp. 738–747. DOI: 10.1086/516845.

Ostrom, Elinor (2010). “Beyond Markets and States: Polycentric Governance of Complex Economic Systems”. In: *American Economic Review* 100.3, pp. 641–672. DOI: 10.1257/aer.100.3.641.

Ott, Edward, Celso Grebogi, and James A. Yorke (1990). “Controlling Chaos”. In: *Physical Review Letters* 64.11, pp. 1196–1199. DOI: 10.1103/PhysRevLett.64.1196.

Parkhurst, Justin (2017). *The Politics of Evidence: From Evidence-Based Policy to the Good Governance of Evidence*. Abingdon: Routledge. ISBN: 9781138570382. URL: https://www.routledge.com/The-Politics-of-Evidence-From-evidence-based-policy-to-the-good-governance/Parkhurst/p/book/9781138570382.

Peter, Fabienne (2023). “Political Legitimacy”. In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta and Uri Nodelman. Winter 2023. Metaphysics Research Lab, Stanford University. URL: https://plato.stanford.edu/archives/win2023/entries/legitimacy/.

Pierson, Paul (1993). “When Effect Becomes Cause: Policy Feedback and Political Change”. In: *World Politics* 45.4, pp. 595–628. DOI: 10.2307/2950710.

— (2000). “Increasing Returns, Path Dependence, and the Study of Politics”. In: *American Political Science Review* 94.2, pp. 251–267. DOI: 10.2307/2586011.

Rancière, Jacques (2013). *The Politics of Aesthetics: The Distribution of the Sensible*. Trans. by Gabriel Rockhill. London: Bloomsbury Academic. ISBN: 9781780936871. URL: https://www.bloomsbury.com/us/politics-of-aesthetics-9781780936871/.

Raphals, Lisa (2016). “Sunzi versus Xunzi: Two Views of Deception and Indirection”. In: *Early China* 39, pp. 185–229. DOI: 10.1017/eac.2016.6.

Robeyns, Ingrid and Morten Fibieger Byskov (2025). “The Capability Approach”. In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta and Uri Nodelman. Summer 2025. Metaphysics Research Lab, Stanford University. URL: https://plato.stanford.edu/archives/sum2025/entries/capability-approach/.

Scheffer, Marten, Jordi Bascompte, William A. Brock, Victor Brovkin, Stephen R. Carpenter, Vasilis Dakos, Hermann Held, Egbert H. van Nes, Max Rietkerk, and George Sugihara (2009). “Early-Warning Signals for Critical Transitions”. In: *Nature* 461, pp. 53–59. DOI: 10.1038/nature08227.

Scheffer, Marten, Steve Carpenter, Jonathan A. Foley, Carl Folke, and Brian Walker (2001). “Catastrophic Shifts in Ecosystems”. In: *Nature* 413, pp. 591–596. DOI: 10.1038/35098000.

Shinbrot, Troy, Edward Ott, Celso Grebogi, and James A. Yorke (1990). “Using Chaos to Direct Trajectories to Targets”. In: *Physical Review Letters* 65.26, pp. 3215–3218. DOI: 10.1103/PhysRevLett.65.3215.

Slingerland, Edward (2003). *Effortless Action: Wu-wei as Conceptual Metaphor and Spiritual Ideal in Early China*. New York: Oxford University Press. ISBN: 9780195138993.

Spiekermann, Kai (2022). “Irreversible Loss”. In: *The Oxford Handbook of Intergenerational Ethics*. Ed. by Stephen M. Gardiner. Oxford University Press, pp. 416–431. DOI: 10.1093/oxfordhb/9780190881931.013.26. URL: https://academic.oup.com/edited-volume/34721/chapter-abstract/386765345.

Tarsney, Christian, Teruji Thomas, and William MacAskill (2024). “Moral Decision-Making Under Uncertainty”. In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta and Uri Nodelman. Metaphysics Research Lab, Stanford University. URL: https://plato.stanford.edu/entries/moral-decision-uncertainty/.

Tömmel, Tatjana and Maurizio Passerin d’Entreves (2025). “Hannah Arendt”. In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta and Uri Nodelman. Spring 2025. Metaphysics Research Lab, Stanford University. URL: https://plato.stanford.edu/archives/spr2025/entries/arendt/.

Valmisa, Mercedes (2025). *All Things Act*. Oxford University Press. ISBN: 9780197812181. DOI: 10.1093/9780197812211.001.0001.

Yue, Feng, Sybille Persson, and David Wasieleski (2023). “Beyond the Western Conceptions with François Jullien: The Process of Leadership Informed by the Chinese Shi as Organizational Propensity”. In: *Management and Organization Review* 19.6, pp. 1050–1070. DOI: 10.1017/mor.2023.17.