

Symmetry, Dynamics, and Emergent Justice: A Field-Theoretic Framework for Economics

Wanhong

huangwanhong.g.official@gmail.com

Preliminary draft, May 5, 2026

Abstract

We develop a unified formal framework connecting welfare economics, algebraic group theory, gauge field theory, and the theory of dynamical systems. The indifference curves of consumer theory are identified as orbits of a one-parameter Lie group acting smoothly on the consumption manifold; this identification is made concrete with explicit computations for Cobb–Douglas, CES, and quasilinear utility functions. Budget constraints are recast as spontaneous symmetry breaking that selects a point from the orbit, and price signals are identified, via the Goldstone–Nambu theorem, with the cotangent directions corresponding to the broken generators. The marginal rate of substitution is shown to equal the ratio of components of the connection one-form on the preference fibre bundle, and its variation along the indifference curve equals the bundle curvature, which is zero for perfect substitutes and generically non-zero for convex preferences. The social path integral is formulated in Euclidean signature, yielding a rigorous stochastic interpretation via the Feynman–Kac formula. We argue that the equilibrium paradigm in economics rests on an ergodicity assumption that is empirically unwarranted for path-dependent social systems; convergence to the Pareto frontier is generically not guaranteed. Building on this analysis, we survey classical ethics (Aristotle, Kant, Mill) and theories of justice (Rawls, Nozick, Sen, Dworkin), develop their formal representations within the dynamical-systems framework, and propose *generative justice*: a functional on trajectory space that evaluates the normative quality of each symmetry-breaking event rather than any terminal state. Open problems are explicitly demarcated throughout.

Keywords. Pareto efficiency; group orbit; spontaneous symmetry breaking; Goldstone theorem; fibre bundle; curvature; Euclidean path integral; ergodicity; generative justice.

MSC 2020. 91B16, 81T13, 37N40, 57R22, 91B14.

JEL. D01, D60, D63, C65, I10.

Contents

1 Introduction

2 Preliminaries

2.1	Quantum Field Theory	3
2.1.1	Lagrangian and Action	3
2.1.2	Euclidean Path Integral and Feynman–Kac	4
2.1.3	Gauge Fields and Curvature	4
2.1.4	Spontaneous Symmetry Breaking	4
2.2	Group Theory	4
2.2.1	Lie Groups and Algebras	4
2.2.2	Orbits, Stabilisers, and the Orbit-Stabiliser Theorem	4
2.2.3	Noether’s Theorem	4
2.3	Dynamical Systems Theory	4
2.3.1	Flows, Fixed Points, and Stability	4
2.3.2	Lyapunov Functions	4
2.3.3	Bifurcation Theory	5
2.3.4	Hamiltonian Systems	5
2.3.5	Ergodic Theory	5
2.3.6	Renormalisation Group as a Dynamical System	5
2.4	Fibre Bundles and Connections	5

3 Economic Concepts

3.1	Consumer Theory	5
3.1.1	Preferences and Utility	5
3.1.2	Revealed Preference Theory	6
3.1.3	Utility Maximisation and Duality	6
3.1.4	Demand Elasticities	6
3.1.5	Consumer Surplus and Welfare Measurement	6
3.1.6	Intertemporal Choice	7
3.1.7	Expected Utility and Risk	7
3.1.8	Engel Curves and Income Effects	7
3.2	Production Theory	7
3.2.1	Production Function and Isoquants	7
3.2.2	Short Run, Long Run, and Cost Structure	7
3.2.3	Profit Function and Hotelling’s Lemma	7
3.2.4	Envelope Theorem	8
3.2.5	Marginal Rate of Transformation and the Complete Pareto Conditions	8

77	3.2.6	Farrell Efficiency Decomposition . . .	8	130	4.9	Pareto Frontier as Boundary Set	18
78	3.3	Market Structures	8	131	4.10	Non-Ergodicity of Social Systems	18
79	3.3.1	Perfect Competition and Deadweight Loss	8	132	4.11	Duality Triangle: Roy, Shephard, and Hotelling as a Single Geometric Structure . . .	18
80				133	4.12	Revealed Preference as Orbit Consistency . .	19
81	3.3.2	Monopoly and the Lerner Index . . .	8	134	4.13	Risk Aversion as Utility Curvature and Stochastic Orbit Measures	19
82	3.3.3	Game Theory: Foundational Concepts . .	9	135	4.14	Intertemporal Choice as Dynamical System; Present Bias as Non-Autonomous Flow . . .	20
83	3.4	General Equilibrium and Welfare Economics . .	9	136	4.15	Homothetic Preferences as Scale-Symmetric Bundles; Engel Curves as Radial Orbits . . .	20
84	3.4.1	Edgeworth Box and Contract Curve . .	9	137	4.16	The Grand Pareto Conditions as Triple Tangency of Manifolds	20
85	3.4.2	Walrasian Equilibrium and Welfare Theorems	9	138	4.17	Short Run vs. Long Run as Constrained vs. Unconstrained Dynamics	21
86				139	4.18	Nash Equilibrium as Fixed Point of Best-Response Correspondence	21
87	3.4.3	Social Welfare Functions	9	140	4.19	Repeated Games as Discrete Dynamical Systems; Folk Theorem as Invariant Set	22
88	3.4.4	Arrow's Impossibility Theorem . . .	10	141	4.20	Nash Bargaining as Geometric Mean Optimization on Orbit Space	22
89	3.4.5	Inequality Measurement	10	142	4.21	Edgeworth Box as Product Bundle; Contract Curve as Tangency Locus	23
90	3.4.6	Externalities and Public Goods . . .	10	143	4.22	Arrow Impossibility as Categorical No-Go; SWF as Aggregation Functor	23
91	3.4.7	Optimal Taxation	10	144	4.23	Lorenz Curve as L^1 Distance; Stochastic Dominance as Partial Order on Distributions . . .	23
92	3.5	Information Economics	11	145	4.24	Optimal Taxation as Variational Problem on Orbit Space	24
93	3.5.1	Adverse Selection and Market Unravelling	11	146	4.25	Coase Theorem as Gauge Invariance of Efficient Outcomes	24
94				147	4.26	Information Asymmetry as Obstruction to Orbit Separation	25
95	3.5.2	Moral Hazard and Principal-Agent Problems	11	148	4.27	Bounded Rationality and Behavioural Concepts as Perturbations of Optimal Orbits . . .	25
96	3.5.3	Signalling	11	149	4.28	Bellman Equation as Value Function Fixed Point; POMDP and Belief-State Geometry . .	26
97	3.5.4	Screening and the Revelation Principle . .	11	150	4.29	Health State Space as Product Manifold; QALY and DALY as Measures	26
98	3.5.5	Supplier-Induced Demand	11	151	4.30	Updated Unified Correspondence Table . . .	27
99				152			
100	3.6	Health Economics	11	153			
101	3.6.1	The Concept of Need in Health . . .	11	154			
102	3.6.2	Grossman Model of Health Demand . .	12	155			
103	3.6.3	Payment Mechanisms and Provider Incentives	12	156			
104				157			
105	3.6.4	Economic Evaluation: Full Typology . .	12	158			
106	3.6.5	Health-Related Quality of Life: EQ-5D . .	12	159			
107	3.6.6	DALY: Disability-Adjusted Life Year . .	12	160			
108	3.6.7	Equity in Health	13	161			
109	3.7	Behavioral Economics	13	162			
110	3.7.1	Bounded Rationality	13	163			
111	3.7.2	Prospect Theory and Loss Aversion . .	13	164			
112	3.7.3	Present Bias and Time Inconsistency . .	13	165			
113	3.7.4	Framing Effects and Reference Dependence	13	166			
114				167			
115	3.8	Dynamic Economics	14	168			
116	3.8.1	Intertemporal Choice and Social Discounting	14	169			
117				170			
118	3.8.2	Dynamic Programming and Bellman Equation	14	171			
119				172			
120	4	Algebraic and Physical Interpretation . .	14	173			
121	4.1	Symmetry Group of a Utility Function: Exact Construction	14	174			
122	4.2	Utility as Conserved Hamiltonian	15	175			
123	4.3	Preference Fibre Bundle and Connection . .	15	176			
124	4.4	Bundle Curvature and Preference Non-Linearity . .	16	177			
125	4.5	Budget Constraints as Symmetry Breaking . .	16	178			
126	4.6	Price as Goldstone Covector	16	179			
127	4.7	Market Equilibrium as a Dynamical Fixed Point . .	17	180			
128	4.8	Social Path Integral in Euclidean Signature . .	17	181			
129							
					5	Ethics, Justice, and Formal Representations . .	27
					5.1	Classical Ethical Frameworks	27
					5.1.1	Virtue Ethics: Aristotle	27
					5.1.2	Kantian Deontology	27
					5.1.3	Utilitarianism	27
					5.1.4	Rawlsian Contractarianism	28
					5.1.5	Libertarianism: Nozick	28
					5.1.6	Capability Approach: Sen	28
					5.1.7	Luck Egalitarianism: Dworkin	28
					5.2	Justice in Economic Theory	28
					5.2.1	The Fundamental Critique of Pareto Optimality	28

182	5.2.2	Social Welfare Functions and Arrow's	
183		Theorem	28
184	5.2.3	Justice Criteria and Welfare Functions	28
185	5.3	Formal Representation of Justice in the Alge-	
186		braic Framework	28
187	5.3.1	Terminal Justice as Fixed-Point Property	28
188	5.3.2	Procedural Justice as Gauge Symmetry	29
189	5.3.3	Capability Justice as Minimum Section	29
190	5.3.4	Generative Justice	29
191	5.3.5	Connection to Biosemiotics and Uni-	
192		versal Grammar	29
193	6	Discussion	29
194	6.1	Implications for Economic Theory	29
195	6.2	Limitations and Open Problems	30
196	7	Related Work	30
197	8	Conclusion	30

198 1 Introduction

199 The mathematical apparatus of neoclassical economics was
 200 assembled during the 1870s by Jevons, Walras, and Edgeworth
 201 through an explicit but largely unreflective borrowing from
 202 classical mechanics [Mirowski, 1989]. Utility was modelled
 203 on potential energy, equilibrium prices on force-balance condi-
 204 tions, and the optimising agent on a particle seeking an energy
 205 minimum. The conceptual consequences of the analogy, and
 206 its limits, were rarely examined.

207 The present paper undertakes a more deliberate and struc-
 208 turally motivated synthesis. We draw on three developments in
 209 20th-century theoretical physics: gauge field theory [Weinberg,
 210 1995], the Goldstone–Nambu theorem on spontaneous sym-
 211 metry breaking [Goldstone, 1961, Nambu and Jona-Lasinio,
 212 1960], and the Euclidean path-integral and Feynman–Kac for-
 213 malism [Feynman and Hibbs, 1965, Simon, 1979]. We also
 214 use the orbit theory of Lie groups [Hall, 2015] and the theory
 215 of connections on fibre bundles [Nakahara, 2003], together
 216 with the modern theory of dynamical systems [Strogatz, 2018,
 217 Guckenheimer and Holmes, 2013].

218 Central claims.

219 (I) The indifference curves of consumer theory are the orbits
 220 of a one-parameter Lie group G_U acting smoothly on the
 221 consumption space X . The utility function U is the invariant
 222 of this action and separates orbits. This correspondence is
 223 exact and requires no analogy.

224 (II) The variation of U along an orbit is zero by definition. Via
 225 the Hamiltonian interpretation of the orbit flow (Claim III),
 226 this is the statement that U is the conserved Hamiltonian,
 227 and Noether's theorem identifies it as the charge of time-
 228 translation symmetry.

229 (III) The orbit flow on X is the Hamiltonian flow of U with
 230 respect to a symplectic structure on X , which for $n = 2$
 231 goods is the standard area form on $\mathbb{R}_+^2$.

232 (IV) Budget constraints break the orbit symmetry G_U , select-
 233 ing a unique point x^* from $\mathcal{I}_c^*$. The price vector p is the
 234 covector $dU(x^*)$ evaluated at the optimum and normalised
 235 by the Lagrange multiplier, playing the role of a Goldstone
 236 covector for the broken symmetry.

237 (V) The marginal rate of substitution equals the ratio of com-
 238 ponents of the connection one-form on the preference fibre
 239 bundle $\pi : X \rightarrow \mathbb{R}$, and the curvature of this connection is
 240 zero if and only if preferences are linearly separable (perfect
 241 substitutes).

242 (VI) The social path integral is well-defined in Euclidean signa-
 243 ture via the Feynman–Kac formula and corresponds to a
 244 stochastic welfare optimisation. The parameter $\hbar_{\text{soc}}$ is an
 245 institutional noise level.

246 (VII) The ergodicity assumption implicit in equilibrium eco-
 247 nomics fails for path-dependent social systems; the Pareto
 248 frontier need not be reachable from given initial conditions.

249 (VIII) Generative justice is a functional on social trajectory space
 250 that evaluates each symmetry-breaking event normatively,
 251 superseding both terminal and procedural conceptions of
 252 justice in non-ergodic systems.

253 Claims (I)–(V) are proved below under clearly stated as-
 254 sumptions. Claims (VI) and (VII) invoke standard results.
 255 Claim (VIII) is a formal proposal whose normative complete-
 256 ness depends on the choice of local scoring function j , an open
 257 problem.

258 The paper is structured as follows. Section 2 provides mathe-
 259 matical and physical background. Section 3 presents economic
 260 concepts with full formulas. Section 4 develops the algebraic
 261 and physical interpretation. Section 5 covers ethics, justice
 262 theory, and their formal representations. Section 6 discusses
 263 implications and limitations. Section 7 surveys related litera-
 264 ture. Section 8 concludes.

265 2 Preliminaries

266 2.1 Quantum Field Theory

267 2.1.1 Lagrangian and Action

268 A classical field theory over a spacetime manifold $\mathcal{M}$ is speci-
 269 fied by a Lagrangian density $\mathcal{L}(\phi, \partial_\mu \phi)$, where $\phi : \mathcal{M} \rightarrow V$ is
 270 a field and ∂_μ denotes spacetime differentiation. The action is

$$S[\phi] = \int_{\mathcal{M}} \mathcal{L}(\phi(x), \partial_\mu \phi(x)) \, d^4x. \quad (1)$$

271 The equations of motion $\delta S = 0$ yield the Euler–Lagrange
 272 equations

$$\frac{\partial \mathcal{L}}{\partial \phi} - \partial_\mu \frac{\partial \mathcal{L}}{\partial (\partial_\mu \phi)} = 0. \quad (2)$$

2.1.2 Euclidean Path Integral and Feynman–Kac

In the Minkowski path integral, amplitudes involve $e^{iS/\hbar}$, which is complex and not directly a probability. The *Euclidean path integral* is obtained by the Wick rotation $t \mapsto -i\tau$ (imaginary time), yielding

$$K(x_f, x_i; T) = \int_{x(0)=x_i}^{x(T)=x_f} \mathcal{D}[x] \exp\left(-\frac{1}{\hbar_{\text{soc}}} S_E[x]\right), \quad (3)$$

where the Euclidean action is

$$S_E[x] = \int_0^T \left[\frac{1}{2} |\dot{x}|^2 + V(x) \right] d\tau. \quad (4)$$

The **Feynman–Kac formula** [Simon, 1979] gives the rigorous interpretation: $K(x_f, x_i; T) = \mathbb{E}^{x_i} \left[e^{-\int_0^T V(x_\tau) d\tau} \mathbf{1}_{x_T=x_f} \right]$, where the expectation is over Brownian paths starting at x_i with diffusion coefficient $\hbar_{\text{soc}}$. This identifies (3) as a genuine probability kernel and connects it to the Fokker–Planck equation

$$\frac{\partial P}{\partial \tau} = \frac{\hbar_{\text{soc}}}{2} \nabla^2 P + \nabla \cdot (P \nabla V). \quad (5)$$

Caveat 2.1. The Minkowski path integral $\int \mathcal{D}[x] e^{iS/\hbar}$ used in some physics-economics translations [Smith and Foley, 2008] is not a probability measure and requires additional regularisation. All social path integrals in this paper use the Euclidean form (3) and are interpreted via the Feynman–Kac formula.

2.1.3 Gauge Fields and Curvature

A *gauge field* is a connection on a principal G -bundle over a base manifold $\mathcal{M}$. In local coordinates, it is a $\mathfrak{g}$ -valued one-form $A = A_\mu dx^\mu$. The *field strength* (curvature) is

$$F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu + [A_\mu, A_\nu]. \quad (6)$$

For an Abelian group $G = U(1)$, the commutator vanishes and $F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu$ is the ordinary electromagnetic tensor. Physical observables are gauge-invariant quantities built from $F_{\mu\nu}$.

2.1.4 Spontaneous Symmetry Breaking

A symmetry is *spontaneously broken* when the action $S[\phi]$ is invariant under a group G but the ground state is not [Nambu and Jona-Lasinio, 1960]. The canonical example is a complex scalar with potential $V(\phi) = -\frac{\mu^2}{2} |\phi|^2 + \frac{\lambda}{4} |\phi|^4$, invariant under $U(1)$, but whose minimum ring $|\phi| = \mu/\sqrt{\lambda}$ is not $U(1)$ -invariant.

Theorem 2.1 (Goldstone–Nambu, 1961). *If a continuous symmetry group G of dimension n is spontaneously broken to a subgroup H of dimension k , then the low-energy spectrum contains at least $n - k$ massless scalar particles (Goldstone bosons). In the absence of gauge fields, the count is exactly $n - k$.*

Proof. Goldstone [1961]; for the general proof, Weinberg [1995], §19.2. $\square$

2.2 Group Theory

2.2.1 Lie Groups and Algebras

Definition 2.1 (Lie group). A *Lie group* is a group G that is also a smooth manifold such that the multiplication $G \times G \rightarrow G$ and inversion $G \rightarrow G$ are smooth maps [Hall, 2015].

The tangent space at the identity, $\mathfrak{g} = T_e G$, with the Lie bracket $[\cdot, \cdot]$ induced by the adjoint action, is the *Lie algebra* of G . The exponential map $\exp : \mathfrak{g} \rightarrow G$ converts Lie algebra elements to one-parameter subgroups.

2.2.2 Orbits, Stabilisers, and the Orbit-Stabiliser Theorem

Definition 2.2. Let G act on a smooth manifold X via $\phi : G \times X \rightarrow X$. The *orbit* of $x \in X$ is $\text{Orb}(x) = G \cdot x$ and the *stabiliser* is $\text{Stab}(x) = \{g \in G : g \cdot x = x\}$.

Theorem 2.2 (Orbit-Stabiliser). $G/\text{Stab}(x) \cong \text{Orb}(x)$ as smooth manifolds when the action is smooth. In particular, $\dim \text{Orb}(x) = \dim G - \dim \text{Stab}(x)$.

Proof. Lee [2013], Theorem 21.18. $\square$

2.2.3 Noether’s Theorem

Theorem 2.3 (Noether, 1918). *Let the action $S[\phi]$ be invariant under a continuous one-parameter family of transformations $\phi \mapsto \phi + \varepsilon \delta\phi + O(\varepsilon^2)$. Then the Noether current $j^\mu = \frac{\partial \mathcal{L}}{\partial(\partial_\mu \phi)} \delta\phi$ satisfies $\partial_\mu j^\mu = 0$ on solutions of the equations of motion, yielding a conserved charge $Q = \int j^0 d^{n-1}x$.*

Proof. Noether [1918]. $\square$

2.3 Dynamical Systems Theory

2.3.1 Flows, Fixed Points, and Stability

A *smooth dynamical system* on a manifold M is specified by a vector field $f : M \rightarrow TM$, generating the ODE $\dot{x} = f(x)$. A *fixed point* satisfies $f(x^*) = 0$. Stability is classified by the eigenvalues of $Df(x^*)$: all eigenvalues with negative real part give an asymptotically stable sink; mixed signs give a saddle; any positive real part gives an unstable fixed point [Strogatz, 2018].

2.3.2 Lyapunov Functions

Definition 2.3. A smooth function $V : U \rightarrow \mathbb{R}$ on a neighbourhood U of a fixed point x^* is a *Lyapunov function* if $V(x^*) = 0$, $V(x) > 0$ for $x \neq x^*$, and $\dot{V}(x) = \nabla V(x) \cdot f(x) \leq 0$.

Theorem 2.4 (Lyapunov Stability). *If a Lyapunov function exists on U , then x^* is Lyapunov-stable. If $\dot{V} < 0$ for $x \neq x^*$, then x^* is asymptotically stable.*

2.3.3 Bifurcation Theory

A *bifurcation* at parameter value μ_c is a qualitative change in the phase portrait of $\dot{x} = f(x; \mu)$ as μ crosses μ_c [Guckenheimer and Holmes, 2013]. Relevant cases:

- *Saddle-node*: two fixed points collide and annihilate; rank $Df = n - 1$ at the bifurcation.
- *Pitchfork*: a symmetric fixed point loses stability and bifurcates into two. This is the dynamical analogue of spontaneous symmetry breaking.
- *Hopf*: a fixed point loses stability, emitting a limit cycle; complex conjugate eigenvalues cross the imaginary axis.

2.3.4 Hamiltonian Systems

A *Hamiltonian system* on a symplectic manifold (M, ω) with Hamiltonian $H : M \rightarrow \mathbb{R}$ satisfies Hamilton's equations:

$$\dot{q}^i = \frac{\partial H}{\partial p_i}, \quad \dot{p}_i = -\frac{\partial H}{\partial q^i}. \quad (7)$$

Key properties: H is conserved along trajectories ($\dot{H} = \{H, H\} = 0$); the flow preserves the symplectic volume form ω^n (Liouville's theorem); and Noether's theorem in the Hamiltonian setting states that each continuous symmetry yields a conserved quantity in involution with H .

2.3.5 Ergodic Theory

Let $(\Omega, \mathcal{F}, \mu, T)$ be a measure-preserving dynamical system.

Definition 2.4. The system is *ergodic* if every T -invariant set has measure 0 or 1.

Theorem 2.5 (Birkhoff Ergodic Theorem). *If the system is ergodic and $f \in L^1(\mu)$, then*

$$\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{k=0}^{N-1} f(T^k \omega) = \int f \, d\mu \quad \mu\text{-a.e.} \quad (8)$$

Proof. Walters [1982], Theorem 1.14. $\square$

Remark 2.1. A system converging to a fixed point x^* concentrates its time-average measure on the singleton $\{x^*\}$, which is a zero-measure set if the state space has positive dimension. Hence asymptotic convergence to a fixed point implies non-ergodicity with respect to any continuous invariant measure. Ergodicity and fixed-point convergence are mutually exclusive properties.

2.3.6 Renormalisation Group as a Dynamical System

The renormalisation group (RG) flow in quantum field theory is a dynamical system on coupling-constant space [Wilson, 1974]:

$$\mu \frac{\partial g_i}{\partial \mu} = \beta_i(g_1, g_2, \dots), \quad (9)$$

where μ is the energy scale. Fixed points of (9) are conformal field theories; their universality classes classify critical phenomena independently of microscopic details. The analogue in macroeconomics is the emergence of robust aggregate regularities (power laws, business cycle patterns) that are independent of individual agent details.

2.4 Fibre Bundles and Connections

Definition 2.5 (Fibre bundle). A *smooth fibre bundle* (E, B, π, F) consists of smooth manifolds E (total space) and B (base), a smooth surjection $\pi : E \rightarrow B$, and a typical fibre F such that each point $b \in B$ has a neighbourhood U with $\pi^{-1}(U) \cong U \times F$ via a diffeomorphism compatible with π [Nakahara, 2003].

Definition 2.6 (Connection). A *connection* on a bundle is a smooth splitting of the tangent bundle $TE = VE \oplus HE$, where $VE = \ker(d\pi)$ is the vertical subbundle. The connection one-form $\omega \in \Omega^1(E, \mathfrak{g})$ satisfies $\omega(v) = v$ for $v \in VE$ and is equivariant under the structure group action. The *curvature* is $\Omega = d\omega + \frac{1}{2}[\omega, \omega] \in \Omega^2(E, \mathfrak{g})$.

For an Abelian structure group ($G = \mathbb{R}$ or $U(1)$), the commutator term vanishes and $\Omega = d\omega$. The curvature $\Omega = 0$ if and only if the connection is flat, *i.e.* the horizontal distribution is integrable (Frobenius theorem).

3 Economic Concepts

Each concept is presented in three parts: the widely accepted natural-language definition with its primary or authoritative source, a discussion of intuition and context, and the formal apparatus. This ordering reflects the methodological principle articulated by Samuelson [1947], p. 3: "Mathematics is a language. The purpose of the mathematical formulation is not to obscure but to clarify." The mathematical model is a precision instrument for reasoning about concepts that precede it.

3.1 Consumer Theory

3.1.1 Preferences and Utility

Accepted definition. Samuelson [1938], p. 61, grounded consumer theory in observed choice: "the individual guinea-pig, by his market behaviour, reveals his preference schedule." Arrow [1951], p. 9: " x is preferred to y if the individual would choose x when y is also available." Mas-Colell *et al.* [1995], p. 9: a utility function U represents $\succsim$ if $x \succsim y \iff U(x) \geq U(y)$.

Discussion. Preferences are prior to utility. Utility is a representation device, not a psychological substance. Pareto [1906] recognised that only the ordering matters (ordinality): any strictly increasing transformation $f \circ U$ represents the same preferences, so utility levels cannot be compared across persons without additional cardinal structure. The axioms of rationality (completeness and transitivity) encode the minimum re-

quirements for consistent choice, but are empirically contested [Kahneman, 2000, Tversky and Kahneman, 1986].

Formal definition. A preference relation $\succsim$ on $X = \mathbb{R}_+^n$ is *rational* iff complete and transitive. Under continuity of $\succsim$, Debreu [1954] shows a continuous representing function U exists, unique up to monotone transformation. Standard special cases: Cobb–Douglas $U = \prod x_i^{\alpha_i}$; CES $U = (\sum \alpha_i x_i^\rho)^{1/\rho}$, $\sigma = 1/(1 - \rho)$; quasilinear $U = x_1 + v(x_2)$, $v'' < 0$; homothetic $U(\lambda x) = \lambda U(x)$.

3.1.2 Revealed Preference Theory

Accepted definition. Samuelson [1938], p. 65: “if at price situation p and income y the individual purchases goods q rather than q' , then we say that q has been selected over q' .” Houthakker [1950] extended this to the *Strong Axiom of Revealed Preference* (SARP), characterising data consistent with utility maximisation.

Discussion. Revealed preference provides a theory of demand without assuming utility; it is purely behavioural and empirically falsifiable. Afriat [1967] proved that a finite set of price–quantity observations is consistent with utility maximisation if and only if it satisfies GARP (Generalised Axiom of Revealed Preference), giving a constructive utility function for any GARP-satisfying data.

Formal definition. x is *directly revealed preferred* to y , written $xR^D y$, if $p \cdot x \leq w$ and $x^*(p, w) = x$. **WARP:** if $xR^D y$ and $x \neq y$, then not $yR^D x$. **SARP:** the transitive closure R of R^D is acyclic. **Afriat’s theorem** [Afriat, 1967]: the following are equivalent: (i) data satisfy GARP; (ii) data are rationalisable by a non-satiated utility function; (iii) there exist $u_k, \lambda_k > 0$ with $u_j \leq u_k + \lambda_k p_k \cdot (x_j - x_k)$ for all k, j .

3.1.3 Utility Maximisation and Duality

Accepted definition. Varian [1992], p. 97: “The consumer’s demand function $x^*(p, w)$ gives the optimal bundle as a function of prices and income.” The dual problem, expenditure minimisation, was formalised by Hicks [1939a], Ch. II.

Formal definitions. *Primal (UMP):* $x^*(p, w) = \arg \max_{p \cdot x \leq w} U(x)$. FOCs: $\nabla U(x^*) = \lambda p$, $p \cdot x^* = w$. *Optimality:* $MRS_{ij}(x^*) = p_i/p_j$. *Dual (EMP):* $h(p, \bar{u}) = \arg \min_{U(x) \geq \bar{u}} p \cdot x$; $e(p, \bar{u}) = p \cdot h(p, \bar{u})$. Shephard’s Lemma: $\partial e / \partial p_i = h_i$.

Slutsky equation [Slutsky, 1915]:

$$\frac{\partial x_i^*}{\partial p_j} = \underbrace{\frac{\partial h_i}{\partial p_j}}_{S_{ij}} - x_j^* \frac{\partial x_i^*}{\partial w}. \quad (10)$$

The Slutsky matrix $S = [S_{ij}]$ is symmetric and negative semidefinite [Samuelson, 1947].

Roy’s Identity [Roy, 1947]:

$$x_i^*(p, w) = -\frac{\partial v / \partial p_i}{\partial v / \partial w}, \quad (11)$$

where $v(p, w) = U(x^*(p, w))$ is the indirect utility function. Roy’s identity connects the indirect utility function directly to Marshallian demand, completing the primal–dual triangle: given v , one recovers both x^* (via Roy) and h (via Shephard), without solving optimisation problems.

Discussion. The duality between UMP and EMP is a fundamental mathematical structure in consumer theory. It mirrors the Legendre transform in physics: the indirect utility function $v(p, w)$ plays the role of the Hamiltonian, and the expenditure function $e(p, \bar{u})$ plays the role of the generating function for canonical transformations. This connection is made precise in Section 4.

3.1.4 Demand Elasticities

Accepted definition. Marshall [1890], Book III, Ch. IV: “the elasticity of demand in a market is great or small according as the amount demanded increases much or little for a given fall in price.”

Formal definition. Own-price: $\varepsilon_{ii} = (\partial x_i^* / \partial p_i)(p_i / x_i^*)$. Cross-price: $\varepsilon_{ij} = (\partial x_i^* / \partial p_j)(p_j / x_i^*)$, $i \neq j$. Income: $\eta_i = (\partial x_i^* / \partial w)(w / x_i^*)$. Aggregation constraints: *Cournot:* $\sum_j s_j \varepsilon_{ij} + s_i = 0$; *Engel:* $\sum_i s_i \eta_i = 1$. Goods classification: necessity $0 < \eta_i < 1$; luxury $\eta_i > 1$; inferior $\eta_i < 0$. Health care demand is notably inelastic ($|\varepsilon| \approx 0.1$ – 0.4 ; Manning *et al.* [1987] RAND experiment).

3.1.5 Consumer Surplus and Welfare Measurement

Accepted definition. Marshall [1890], Book III, Ch. VI: “the excess of the price which [a consumer] would be willing to pay rather than go without the thing, over that which he actually does pay.”

Formal definition. Marshallian CS: $CS(p_i) = \int_{p_i}^{\bar{p}} x_i^* d\bar{p}_i$. Exact measures: $CV = e(p^0, u^0) - e(p^1, u^0)$ (compensating variation: income adjustment to restore original utility); $EV = e(p^0, u^1) - e(p^1, u^1)$ (equivalent variation: income change equivalent to the price change). Willig [1976]: $|CS - CV|/CV < \eta_i |\Delta CS|/(2w)$, typically small. For quasilinear utility: $CS = CV = EV$ exactly.

3.1.6 Intertemporal Choice

Accepted definition. Fisher [1930], Ch. IV, modelled saving and borrowing as a choice between present and future consumption. Samuelson [1937] introduced the discounted utility model as its formal expression.

Discussion. Intertemporal preferences determine saving behaviour, investment in health, and responses to health policy. The social discount rate determines the weight given to future generations' welfare, making it central to intergenerational justice [Stern, 2007, Nordhaus, 2008].

Formal definition. Standard model: $U = \sum_{t=0}^T \beta^t u(c_t)$, $\beta = 1/(1+\rho) \in (0, 1)$, ρ the pure rate of time preference. The Euler equation:

$$u'(c_t) = \beta(1 + r_{t+1}) u'(c_{t+1}), \quad (12)$$

characterises the optimal intertemporal consumption path.

Hyperbolic discounting [Laibson, 1997]: $U = u(c_0) + \beta \sum_{t=1}^T \delta^t u(c_t)$, $\beta < 1$, $\delta \approx 1$. The present-bias parameter β generates time-inconsistency: plans made today for tomorrow are not followed through. This is empirically well-documented in health behaviour (smoking, exercise, medication adherence).

3.1.7 Expected Utility and Risk

Accepted definition. von Neumann and Morgenstern [1944], Ch. 3, proved that under four axioms (completeness, transitivity, continuity, independence), a rational agent's preferences over lotteries are represented by expected utility: $U(L) = \sum_s \pi_s u(x_s)$. Arrow [1963], p. 946: "most people would prefer a certain income of \$1,000 to an uncertain income with the same expected value." This is risk aversion.

Formal definition. vNM utility: $EU(L) = \sum_s \pi_s u(x_s)$ for lottery $L = (x_s, \pi_s)_s$. The *certainty equivalent* $CE(L)$: $u(CE) = EU(L)$. *Risk premium*: $\rho = \mathbb{E}[x] - CE(L) \geq 0$ for risk-averse agents. *Arrow-Pratt measure of absolute risk aversion* [Arrow, 1965, Pratt, 1964]:

$$r_A(x) = -\frac{u''(x)}{u'(x)} \geq 0 \quad (\text{risk aversion}). \quad (13)$$

Relative risk aversion: $r_R(x) = x r_A(x)$. CRRA utility: $u(x) = x^{1-\gamma}/(1-\gamma)$, $r_R = \gamma$ constant. Health insurance demand increases with r_A ; the optimal coinsurance rate decreases with risk aversion and increases with moral hazard [Arrow, 1963].

Prospect Theory. Kahneman and Tversky [1979] showed that observed choices under risk systematically violate expected utility. Prospect theory replaces $u(x)$ with a value function $v(x - r)$ defined relative to a reference point r , exhibiting loss

aversion ($v'(x - r) < v'(-(x - r))$ for $x > r$, loss aversion coefficient $\lambda \approx 2.25$) and diminishing sensitivity. Probabilities are weighted by $w(\pi) \neq \pi$, overweighting small probabilities. Prospect theory has applications in health insurance design, patient responses to out-of-pocket costs, and policy framing effects.

3.1.8 Engel Curves and Income Effects

Accepted definition. Engel [1857], p. 28, noted that "the poorer a family, the greater the proportion of total expenditure that must be devoted to food." The Engel curve describes how demand for a good changes with income, holding prices fixed.

Formal definition. The Engel curve for good i is the function $w \mapsto x_i^*(p, w)$ at fixed p . Its slope is $\partial x_i^*/\partial w = x_i^* \eta_i/w$. Working's law [Working, 1943]: budget share $s_i = p_i x_i^*/w$ is approximately linear in $\ln w$. Engel curves are central to measuring living standards and poverty, and to constructing price indices.

3.2 Production Theory

3.2.1 Production Function and Isoquants

Accepted definition. Varian [1992], p. 1: "The production function $f(x)$ gives the maximum amount of output that can be produced using the input bundle x ." Mas-Colell *et al.* [1995], p. 128: the production function "gives, for each input vector z , the maximum level of output that can be produced."

Formal definition. $F : \mathbb{R}_+^m \rightarrow \mathbb{R}_+$, monotone, quasi-concave. Isoquant: $Q_y = \{z : F(z) = y\}$. $MRTS_{ij} = (\partial F/\partial z_i)/(\partial F/\partial z_j)$ is the slope of the isoquant.

3.2.2 Short Run, Long Run, and Cost Structure

Accepted definition. Marshall [1890], Book V, Ch. V: in the short run, some factors are fixed; in the long run, all factors are variable. "The short period is too short to allow of any considerable change in the supply of the more specialised of the [factors]", while "in the long run all investments of capital and effort in providing the material plant [adjust]."

Formal definition. Short-run cost: $c^{SR}(w, y; \bar{z}_k)$ fixes inputs $\bar{z}_k$. Long-run cost: $c^{LR}(w, y) = \min_z w \cdot z$ s.t. $F(z) \geq y$. Always $c^{SR} \geq c^{LR}$ (envelope result). Standard cost components: $TC(y) = FC + VC(y)$; $MC(y) = \partial TC/\partial y$; $AC(y) = TC(y)/y$. Minimum efficient scale (MES): y^{**} where $AC'(y^{**}) = 0$, equivalently $MC = AC$.

3.2.3 Profit Function and Hotelling's Lemma

Accepted definition. Varian [1992], p. 43: "The profit function $\pi(p, w)$ gives the maximum profit as a function of the output price p and the input prices w ."

Formal definition. $\pi(p_y, w_z) = \max_{y,z} [p_y y - w_z \cdot z]$ s.t. $F(z) \geq y$. The profit function is convex and homogeneous of degree 1 in (p_y, w_z) . *Hotelling's Lemma* [Hotelling, 1932]:

$$\frac{\partial \pi}{\partial p_y} = y^*(p_y, w_z), \quad \frac{\partial \pi}{\partial w_{z,i}} = -z_i^*(p_y, w_z). \quad (14)$$

Hotelling's Lemma is the production analogue of Shephard's Lemma and Roy's Identity. Together, the three duality relations (Roy, Shephard, Hotelling) form a complete system connecting value functions to optimal choices via differentiation, without re-solving optimisation.

3.2.4 Envelope Theorem

Accepted definition. Milgrom and Segal [1999]: the envelope theorem states that, at an optimum, the derivative of the optimal value with respect to a parameter equals the partial derivative of the objective with respect to that parameter (holding choice variables fixed at their optimal values). It unifies all duality results in economics.

Formal definition. Let $V(\alpha) = \max_x f(x, \alpha)$ s.t. $g(x, \alpha) \leq 0$. Then:

$$\frac{dV}{d\alpha} = \left. \frac{\partial \mathcal{L}}{\partial \alpha} \right|_{x=x^*(\alpha)}, \quad (15)$$

where $\mathcal{L} = f - \lambda g$ is the Lagrangian. This gives Shephard's Lemma (differentiating e w.r.t. p), Roy's Identity (differentiating v w.r.t. p), and Hotelling's Lemma (differentiating π w.r.t. p_y) as special cases.

3.2.5 Marginal Rate of Transformation and the Complete Pareto Conditions

Accepted definition. Samuelson [1947], p. 234: the *marginal rate of transformation* (MRT) between two goods is "the rate at which the economy can transform one good into another by reallocation of resources." It is the slope of the *production possibility frontier* (PPF). Mas-Colell *et al.* [1995], p. 565: "MRT₁₂ = MC₁/MC₂" at any efficient production point.

Discussion. The complete characterisation of Pareto efficiency in a production economy requires three conditions to hold simultaneously [Samuelson, 1947, Mas-Colell *et al.*, 1995]:

1. *Exchange efficiency*: all consumers have the same MRS for every pair of goods.
2. *Production efficiency*: all firms have the same MRTS for every pair of inputs.
3. *Product mix efficiency*: the common MRS equals the MRT.

Formal definition. For a two-good economy with PPF $F(y_1, y_2) = 0$:

$$\text{MRT}_{12} = - \left. \frac{dy_2}{dy_1} \right|_{F=0} = \frac{MC_1}{MC_2}. \quad (16)$$

The *grand Pareto conditions*:

$$\underbrace{\text{MRS}_{12}^i = \text{MRS}_{12}^j}_{\text{exchange}} = \underbrace{\text{MRT}_{12}}_{\text{product mix}}, \quad \underbrace{\text{MRTS}_{12}^k = \text{MRTS}_{12}^l}_{\text{production}} \quad \forall i, j, k, l. \quad (17)$$

Competitive equilibrium achieves all three conditions via the price system: the common consumer MRS equals p_1/p_2 (from optimality condition (??)); firms with the same input prices achieve the same MRTS; and the product mix condition $\text{MRS} = \text{MRT} = p_1/p_2$ follows from $p_i = MC_i$ under competition. This is the formal content of the First Welfare Theorem.

3.2.6 Farrell Efficiency Decomposition

Accepted definition. Farrell [1957], p. 254: "the productive efficiency of a firm consists of two components: technical efficiency, which reflects the ability of a firm to obtain maximal output from a given set of inputs, and allocative efficiency, which reflects the ability of a firm to use the inputs in optimal proportions."

Formal definition. Technical efficiency (TE): $\text{TE} = F(z)/y_{\max}(z) \in [0, 1]$. Scale efficiency (SE): ratio of productivity at current scale to productivity at MES. Economic efficiency: $\text{EE} = \text{TE} \times \text{AE}$, ratio of minimum feasible cost to actual cost. DEA (Data Envelopment Analysis) [Charnes *et al.*, 1978] estimates the frontier non-parametrically from observed input-output data, providing an empirical implementation of Farrell efficiency for health systems and hospitals.

3.3 Market Structures

3.3.1 Perfect Competition and Deadweight Loss

Accepted definition. Marshall [1890], Book V, Ch. I: competitive markets feature many price-taking buyers and sellers. Mas-Colell *et al.* [1995], p. 315: "a market is perfectly competitive when all agents are price-takers."

Formal definition. Each firm sets $p = MC(y)$. Long-run equilibrium: $p = MC = AC$ (zero profit, free entry). Deadweight loss from any market distortion: $\text{DWL} = \text{CS}^* + \text{PS}^* - \text{CS} - \text{PS} - T$ (starred: competitive benchmark, T : tax revenue).

3.3.2 Monopoly and the Lerner Index

Accepted definition. Tirole [1988], p. 65: "A monopolist is a firm that is the only seller on a market and, as such, faces the entire market demand." Lerner [1934]: the degree of monopoly power is measured by the price-cost margin.

Formal definition. $MR(y) = p(1 + 1/\varepsilon_{yy})$. Monopoly optimum: $MR = MC$. Lerner index: $L = (p - MC)/p = -1/\varepsilon_{yy}$. $L = 0$: perfect competition. $L \rightarrow 1$: extreme monopoly power.

3.3.3 Game Theory: Foundational Concepts

Accepted definition. von Neumann [1928] founded game theory; Nash [1950], p. 48: an equilibrium is a strategy profile from which no player can profitably deviate. Osborne and Rubinstein [1994], p. 14: “a Nash equilibrium is an action profile a^* such that no player can do better by choosing differently, given others’ actions.”

Dominant strategy (Luce and Raiffa [1957], p. 69): a strategy that is optimal regardless of others’ actions. Dominant strategy equilibrium is the strongest solution concept; Nash equilibrium is weaker.

Prisoner’s Dilemma [Tucker, 1950]: two players, each choosing between Cooperate (C) and Defect (D). Payoff matrix (row player):

	C	D
C	(3, 3)	(0, 5)
D	(5, 0)	(1, 1)

D is a dominant strategy for both; the unique Nash equilibrium (D, D) yields payoff (1, 1), Pareto-dominated by (C, C) . Individual rationality produces collective inefficiency. The prisoner’s dilemma is the canonical model of externalities, public goods, arms races, and antibiotic overuse in healthcare.

Formal definitions. *Nash equilibrium*: $s^* \in \prod_i S_i$ with $u_i(s_i^*, s_{-i}^*) \geq u_i(s_i, s_{-i}^*)$ for all s_i, i .

Subgame perfect equilibrium (SPE) [Selten, 1965]: in an extensive-form game, a Nash equilibrium in every subgame. SPE rules out non-credible threats. Backward induction finds SPE in finite games of perfect information.

Repeated games and Folk Theorem [Friedman, 1971]: in the infinitely repeated prisoner’s dilemma with discount factor δ close to 1, any feasible payoff strictly above the Nash equilibrium payoff can be sustained as a subgame perfect equilibrium (via trigger strategies). Formally: payoff vector v is supportable iff $v_i \geq \underline{v}_i$ (minimax) for all i and v is feasible. This underlies the stability of reputation, implicit contracts, and professional norms in medicine.

Nash bargaining solution [Nash, 1950b]: for a bargaining problem (F, d) where F is the feasible payoff set and d the disagreement point, the unique solution satisfying symmetry, Pareto optimality, scale invariance, and independence of irrelevant alternatives is:

$$(u_1^*, u_2^*) = \arg \max_{(u_1, u_2) \in F, u_i \geq d_i} (u_1 - d_1)(u_2 - d_2). \quad (18)$$

This applies to price negotiations between hospitals and insurers, physician-patient treatment decisions, and wage bargaining in healthcare.

3.4 General Equilibrium and Welfare Economics

3.4.1 Edgeworth Box and Contract Curve

Accepted definition. Edgeworth [1881], p. 28, introduced the box diagram for visualising exchange between two agents. Pareto [1906] identified the contract curve as the locus of all Pareto-optimal exchange allocations. Varian [1992], p. 517: “the contract curve is the set of Pareto-efficient allocations in the Edgeworth box.”

Discussion. In a two-person (A, B) , two-good (x_1, x_2) exchange economy, the Edgeworth box has dimensions $\omega_1^A + \omega_1^B$ by $\omega_2^A + \omega_2^B$. Every point in the box represents a feasible allocation. An allocation is Pareto optimal iff the indifference curves of A and B are tangent, i.e. $MRS_{12}^A = MRS_{12}^B$. The set of such tangency points is the *contract curve* (or *Pareto set*). Walrasian equilibria lie on both the contract curve and the *budget line* through the endowment point, as guaranteed by the First Welfare Theorem.

Formal definition. Allocation (x^A, x^B) with $x^A + x^B = \omega^A + \omega^B$ is Pareto optimal iff: $MRS_{12}^A(x^A) = MRS_{12}^B(x^B)$. The contract curve $\mathcal{C} = \{(x^A, x^B) : \nabla U^A(x^A) = \mu \nabla U^B(x^B), \mu > 0, x^A + x^B = \omega\}$ for some scalar μ .

3.4.2 Walrasian Equilibrium and Welfare Theorems

Accepted definition. Walras [1874] defined equilibrium as the price vector at which all markets clear simultaneously. Arrow and Hahn [1971], p. 1: “a price system and choice vectors such that all markets simultaneously clear.”

Formal definition. Excess demand: $z_i(p) = \sum_j x_i^j(p, p \cdot \omega^j) - \sum_j \omega_i^j$. Equilibrium: $z(p^*) = 0$. Walras’s law: $p \cdot z(p) \equiv 0$.

Theorem 3.1 (Arrow–Debreu, 1954). *Under convexity and continuity, a Walrasian equilibrium p^* exists (proved via Kakutani’s fixed-point theorem [Kakutani, 1941]).*

Theorem 3.2 (First Welfare Theorem). *Every Walrasian equilibrium is Pareto optimal (under local non-satiation). Proof: if (x^{j*}) is not Pareto optimal, a dominating allocation (x^j) exists with $p^* \cdot x^j \geq p^* \cdot x^{j*}$ for all j (by revealed preference), summing to a contradiction of market clearing.*

Theorem 3.3 (Second Welfare Theorem). *Every Pareto-optimal allocation can be decentralised as a Walrasian equilibrium after lump-sum redistribution [Debreu, 1954b].*

3.4.3 Social Welfare Functions

Accepted definition. Bergson [1938], p. 310: a social welfare function is “any function of the quantities of commodities consumed and of the amounts of the various kinds of work

done by all of the individuals” that allows social comparisons. Samuelson [1947] developed this into a systematic framework.

Discussion. A Bergson–Samuelson SWF $W : \mathbb{R}^m \rightarrow \mathbb{R}$ maps the vector of individual utilities $(u_1, \dots, u_m)$ to a social welfare indicator. Specific functional forms encode different ethical commitments:

Form	$W(u_1, \dots, u_m)$	Ethic
Utilitarian	$\sum_i u_i$	Bentham [1789]
Prioritarian	$\sum_i \phi(u_i), \phi'' < 0$	Parfit [1984]
Rawlsian	$\min_i u_i$	Rawls [1971]
Nash	$\prod_i u_i$	Nash [1950b]
CES	$(\sum_i u_i^\rho)^{1/\rho}$	General

As $\rho \rightarrow -\infty$, CES converges to Rawlsian maximin. As $\rho \rightarrow 1$, CES converges to utilitarian.

Formal definition. A SWF $W : \mathbb{R}^m \rightarrow \mathbb{R}$ satisfies: (i) *Pareto*: $u_i \geq u'_i$ for all i , with strict inequality for some, implies $W(u) > W(u')$. (ii) *Symmetry*: W is symmetric in its arguments (anonymous). (iii) *Concavity*: W is quasi-concave, implementing aversion to inequality.

3.4.4 Arrow’s Impossibility Theorem

Accepted definition. Arrow [1951], p. 59: “If we exclude the possibility of interpersonal comparisons of utility, then the only methods of passing from individual tastes to social preferences which will be satisfactory and which will be defined for a wide range of sets of individual orderings are either imposed or dictatorial.”

Formal statement. No social welfare function $F : \mathcal{R}^m \rightarrow \mathcal{R}$ satisfies simultaneously: (P) Pareto efficiency; (IIA) independence of irrelevant alternatives; (ND) non-dictatorship; (U) unrestricted domain.

Condorcet Paradox. Condorcet [1785] showed that majority voting can produce cyclic social preferences even when individual preferences are transitive. Example: three voters rank $\{a, b, c\}$ as $a \succ b \succ c$; $b \succ c \succ a$; $c \succ a \succ b$. Majority prefers a to b , b to c , but c to a : a cycle. The Condorcet paradox is a concrete instantiation of Arrow’s theorem showing that aggregation of individual rational preferences can yield collective irrationality.

3.4.5 Inequality Measurement

Accepted definition. Lorenz [1905], p. 209, introduced the Lorenz curve to describe income inequality: “the percentage of population against the percentage of income received.” Gini [1912] defined the Gini coefficient as a scalar summary of inequality.

Formal definition. The *Lorenz curve* $L : [0, 1] \rightarrow [0, 1]$ plots the cumulative income share of the bottom p fraction of the population: $L(p) = F^{-1}(\int_0^p F^{-1}(q) dq) / \mu$, where F is the income CDF and μ the mean income. $L(p) \leq p$ always (“bow-shaped below the diagonal”). *Gini coefficient*:

$$G = 1 - 2 \int_0^1 L(p) dp \in [0, 1]. \quad (19)$$

$G = 0$: perfect equality; $G = 1$: perfect inequality. *Atkinson index* [Atkinson, 1970]: $A_\varepsilon = 1 - (\mu_e / \mu)$, where μ_e is the equally distributed equivalent income, embedding explicit inequality aversion ε .

Stochastic dominance: distribution F *first-order stochastically dominates* G if $F(x) \leq G(x)$ for all x , i.e. higher income at every percentile. Second-order stochastic dominance requires equal means and Lorenz-curve dominance, an inequality ordering consistent with all concave utility functions.

3.4.6 Externalities and Public Goods

Accepted definition. Pigou [1920], Part II, Ch. IX: externality is a divergence between private and social net product. Samuelson [1954]: a pure public good is non-excludable and non-rival.

Pigouvian tax and Coase theorem. *Pigouvian tax*: $t^* = \text{MEC}(y^*)$ at the social optimum [Pigou, 1920]. *Coase theorem* [Coase, 1960], p. 15: “With perfect information and zero transaction costs, the efficient outcome will be achieved regardless of the assignment of property rights.” Private bargaining internalises externalities without government intervention. The Coase theorem and the Pigouvian approach represent two poles: markets can self-correct (Coase) versus government intervention is necessary (Pigou). Transaction costs, information asymmetries, and the scale of the externality determine which approach is appropriate.

Formal definition. Pigouvian optimum: $t^* = \text{MEC}(y^*)$. Samuelson condition for public good G : $\sum_i \text{MRS}_{i,G,x_i} = \text{MC}(G)$.

3.4.7 Optimal Taxation

Accepted definition. Ramsey [1927], p. 47: taxes should be levied “in inverse proportion to the elasticities of demand.” Mirrlees [1971] extended this to optimal nonlinear income taxation under asymmetric information.

Discussion. Optimal taxation balances two objectives: raising revenue (efficiency) and redistributing income (equity). The Ramsey rule minimises deadweight loss for a given revenue requirement; the Mirrlees framework trades off redistribution against incentive costs, producing a progressive but non-confiscatory schedule.

Formal definition. *Ramsey rule:* to raise required revenue R with minimum aggregate DWL, set commodity tax rates $\{t_i\}$ satisfying:

$$\sum_j t_j \frac{\partial x_j^H}{\partial p_i} = -\mu x_i^H \quad \forall i, \quad (20)$$

where x^H is Hicksian demand and μ is the Lagrange multiplier on the revenue constraint. The *inverse elasticity rule*: for independent goods, $t_i/p_i \propto 1/|\varepsilon_{ii}^H|$.

Mirrlees optimal income tax: individuals differ in unobservable skill $n \sim F(n)$. The planner maximises $\int W(u(n)) dF(n)$ subject to incentive compatibility and budget balance. The top marginal rate is determined by the interplay between the social welfare weight and the skill distribution tail.

3.5 Information Economics

3.5.1 Adverse Selection and Market Unravelling

Accepted definition. Akerlof [1970], p. 488: in markets with asymmetric quality information, “bad cars drive out good cars.” Only below-average quality is offered at any price buyers will pay, potentially causing market collapse.

Formal definition. Seller of quality $\theta \sim F$ offers at price p . Only types $\theta \leq \theta^*(p)$ sell; buyer’s expected quality $\mathbb{E}[\theta|\theta \leq \theta^*]$ falls as p falls. If the buyer’s valuation of expected quality falls below p : market unravels. *Rothschild–Stiglitz equilibrium* [Rothschild and Stiglitz, 1976]: in a competitive insurance market with adverse selection, a separating equilibrium exists only if high-risk types receive full insurance and low-risk types receive partial coverage. A pooling equilibrium does not generally exist.

3.5.2 Moral Hazard and Principal-Agent Problems

Accepted definition. Arrow [1963], p. 961: insurance “might itself change incentives.” Holmstrom [1979]: the principal-agent problem arises when effort is unobservable.

Formal definition. Full model: agent chooses effort e ; outcome $\theta \sim F(\cdot|e)$; contract $w(\theta)$. (IR): $\mathbb{E}[u(w) - c(e)|e] \geq \bar{u}$. (IC): $e \in \arg \max_{e'} \mathbb{E}[u(w) - c(e')|e']$. First-best (observable e): full insurance $w = \text{const}$; second-best (unobservable e): $w(\theta)$ must vary with θ to provide incentives, exposing risk-averse agent to risk. *Borch rule*: optimal risk-sharing requires equal marginal utilities of income across states; moral hazard prevents full attainment.

3.5.3 Signalling

Accepted definition. Spence [1973], p. 355: in a competitive labour market with asymmetric information, “education can function as a signal” if it is cheaper for high-ability workers to acquire, separating them from low-ability workers.

Formal definition. Two types $\theta \in \{\theta_L, \theta_H\}$, $\theta_H > \theta_L$, with private information. Signalling cost $c(e, \theta)$, $\partial^2 c / \partial e \partial \theta < 0$ (*single-crossing property*). In a separating equilibrium, high types choose e_H^* satisfying: (i) $\theta_H - c(e_H^*, \theta_H) \geq \theta_L$ (no mimicking); (ii) $\theta_L - c(e_L^*, \theta_L) \leq \theta_L$ (IR for low type). Education as signalling without human capital accumulation is a pure welfare loss. In healthcare: hospital accreditation, physician board certification, and drug regulation function as quality signals.

3.5.4 Screening and the Revelation Principle

Accepted definition. Screening is the uninformed party’s design of a menu of contracts to elicit private information. Myerson [1979]: the *Revelation Principle* states that any equilibrium outcome of any mechanism can be replicated by an incentive-compatible direct revelation mechanism.

Formal definition. *Direct mechanism*: each agent reports type $\hat{\theta}$; allocation $q(\hat{\theta})$, payment $t(\hat{\theta})$. *Incentive compatibility*: $u(q(\theta), t(\theta); \theta) \geq u(q(\hat{\theta}), t(\hat{\theta}); \theta)$ for all $\hat{\theta}$. The revelation principle reduces mechanism design to finding the optimal IC and IR direct mechanism. It underlies optimal health insurance contract design, drug pricing, and hospital payment scheme reform.

3.5.5 Supplier-Induced Demand

Accepted definition. Evans [1974]: physicians exploit information asymmetry to generate demand beyond the patient’s informed optimum. McGuire [2000], p. 462: “Physicians act as the patient’s agent. As agent, they may recommend more or less than the patient would want if fully informed.”

Formal definition. Let D^* be informed demand and D observed demand; $SID = D - D^* > 0$. Degree of agency: $\alpha \in [0, 1]$ measures alignment between physician objective and patient welfare. Physician utility: $U_{Ph} = \pi(q) + \alpha U_P(q)$, where q is quantity, π profit. Perfect agency: $\alpha = 1$, no SID. Complete selfishness: $\alpha = 0$, maximum SID. Fee-for-service payment systems generate stronger SID incentives than capitation [McGuire, 2000].

3.6 Health Economics

3.6.1 The Concept of Need in Health

Accepted definition. Bradshaw [1972] identified four types of need:

1. *Felt need*: what people say they need; subjective and not necessarily expert.
2. *Expressed need*: felt need converted into demand by seeking services.
3. *Normative need*: what an expert or professional defines as need; may differ from felt need.

- 943 4. *Comparative need*: need identified by comparison with a
944 reference group.

945 **Discussion.** The distinction between need and demand is
946 fundamental in health economics and health policy. Healthcare
947 is special: consumers (patients) often cannot assess their own
948 need (information asymmetry), the consequence of unmet need
949 (illness or death) is severe, and the distribution of need is
950 unequal across the population. Phelps [2017], p. 6: “In health
951 care, ‘need’ reflects an assessment by some person or group
952 that someone requires medical care,” distinct from effective
953 demand. This distinction motivates universal health coverage
954 as a policy goal: the goal is to meet normative need, not just
955 effective demand.

956 3.6.2 Grossman Model of Health Demand

957 **Accepted definition.** Grossman [1972], p. 223: “health can
958 be viewed as a durable capital stock that produces an output of
959 healthy time. Individuals inherit an initial stock of health that
960 depreciates with age and can be increased by investment.”

961 **Formal model.** $H_{t+1} = H_t(1 - \delta_t) + I(m_t, TH_t)$, δ_t in-
962 creasing in age. Optimal investment equates the marginal cost
963 of health investment to the marginal return from an additional
964 unit of health stock (in terms of productive healthy time and
965 utility). Health demand is derived from the demand for health
966 capital; healthcare demand is derived from health demand. This
967 explains why demand for medical care increases with wages
968 (value of time), decreases with age (rising depreciation), and
969 increases with education (higher investment efficiency).

970 3.6.3 Payment Mechanisms and Provider Incentives

971 **Accepted definition.** Ellis and McGuire [1986]: alterna-
972 tive payment systems create different incentive structures for
973 providers. McGuire [2000], p. 502: “The payment method is
974 the key policy instrument for influencing physician behavior.”

975 **Discussion.** The principal mechanism design problem in
976 health systems is designing payments that align provider incen-
977 tives with patient welfare and system efficiency.

Payment	Incentive	Risk bearer
Fee-for-service (FFS)	Volume maximisation	Payer
DRG/case-based	Cost minimisation	Provider
Capitation	Population health	Provider
Pay-for-performance	Quality indicators	Provider
Bundled payment	Episode efficiency	Provider

979 FFS generates the strongest SID incentive; capitation the
980 weakest. DRG payment (introduced in US Medicare 1983)
981 reduced average length of stay substantially [Ellis and McGuire,

1986]. The optimal payment is the solution to a principal-
983 agent problem: balance provider risk-bearing against incentive
984 intensity.

985 3.6.4 Economic Evaluation: Full Typology

986 **Accepted definition.** Drummond *et al.* [2015], p. 6: “Eco-
987 nomic evaluation is the comparative analysis of alternative
988 courses of action in terms of both their costs and conse-
989 quences.”

990 Full typology.

- 991 1. *Cost minimisation analysis (CMA)*: used when outcomes
992 are demonstrably equivalent; compare costs only.
- 993 2. *Cost-effectiveness analysis (CEA)*: outcomes measured in
994 natural units (*e.g.* life years gained, infections prevented).
995 $ICER_{CEA} = \Delta C / \Delta E_{\text{natural}}$.
- 996 3. *Cost-utility analysis (CUA)*: outcomes measured in
997 QALYs, enabling cross-disease comparison. $ICER_{CUA}$
998 $= \Delta C / \Delta QALY$.
- 999 4. *Cost-benefit analysis (CBA)*: all outcomes monetised, allow-
1000 ing comparison across sectors. Net benefit $= \lambda \Delta E - \Delta C$,
1001 where λ is the willingness-to-pay per unit of outcome.

1002 **Formal definition.** QALY: $QALY = q \times t$, $q \in [0, 1]$ qual-
1003 ity weight. ICER: $\Delta C / \Delta E$. Net monetary benefit (NMB):
1004 $NMB = \lambda \cdot \Delta QALY - \Delta C$; adopt if $NMB > 0$ [Stinnett
1005 and Mullahy, 1998]. NICE threshold: $\lambda \in \text{GBP } 20,000\text{--}$
1006 $30,000/\text{QALY}$ [NICE, 2013].

1007 3.6.5 Health-Related Quality of Life: EQ-5D

1008 **Accepted definition.** EuroQol Group [1990]: EQ-5D is a
1009 standardised instrument for measuring health-related quality
1010 of life. It covers five dimensions: mobility, self-care, usual
1011 activities, pain/discomfort, anxiety/depression; each at 3 (EQ-
1012 5D-3L) or 5 (EQ-5D-5L) severity levels.

1013 **Discussion.** EQ-5D is the most widely used instrument for
1014 generating QALY utility weights in health technology assess-
1015 ment. For the 5L version, $5^5 = 3,125$ health states are possi-
1016 ble; utility values for each state are derived from population-
1017 based valuation studies (time trade-off or discrete choice ex-
1018 periments). The UK EQ-5D-5L value set [Devlin *et al.*, 2018]
1019 maps health states to utility values $\in [-0.285, 1]$, where states
1020 worse than death receive negative weights.

1021 3.6.6 DALY: Disability-Adjusted Life Year

1022 **Accepted definition.** Murray and Lopez [1994], p. 429: “The
1023 DALY extends the concept of potential years of life lost due to
1024 premature death to include equivalent years of healthy life lost

by virtue of being in states of poor health.” DALYs measure the burden of disease rather than the benefit of intervention (as QALYs do).

Formal definition. DALY = YLL + YLD, where YLL (years of life lost) = $N \times L$ (N deaths, L standard life expectancy at age of death) and YLD (years lived with disability) = $I \times DW \times L$ (I incident cases, $DW \in [0, 1]$ disability weight, L average duration). DALYs averted by an intervention are the effectiveness measure in CUA from the burden-of-disease perspective. One DALY $\approx$ one lost year of healthy life; $1 - DW \approx q$ (QALY weight) but the two measures embed different value judgments about age weighting, social discounting, and the relationship between disability and quality of life.

3.6.7 Equity in Health

Accepted definition. Mooney [1987]: equity requires equal access for equal need, equal utilisation for equal need, or equal health outcomes. Williams [1997], p. 117: the *fair innings* argument: “there is a fair innings which is represented by a reasonable span of health, and that everyone is entitled to such a fair innings.” Those who have not yet had their fair innings have a stronger claim on healthcare resources.

Discussion. Efficiency and equity are frequently in tension. Maximising aggregate QALYs concentrates resources on those with greatest benefit potential (typically younger, healthier patients), systematically disadvantaging the elderly, disabled, and chronically ill. The rule of rescue [Jonsen, 1986] reflects a deontological pull toward identifiable individuals in immediate peril, overriding utilitarian calculations. Addressing equity in health resource allocation is a justice problem that no purely technical framework resolves; it requires the normative analysis developed in Section 5.

3.7 Behavioral Economics

3.7.1 Bounded Rationality

Accepted definition. Simon [1955], p. 101: “the capacity of the human mind for formulating and solving complex problems is very small compared with the size of the problems whose solution is required for objectively rational behavior.” Agents *satisfice*: they search until they find an option meeting an aspiration level, rather than optimising.

Discussion. Bounded rationality challenges the neoclassical model at its foundation. In healthcare: patients use heuristics to evaluate treatment options; physicians apply clinical rules of thumb that deviate from Bayesian updating; policy makers engage in framing-dependent choices. Bounded rationality has implications for health education, nudge-based policy design, and the interpretation of “expressed demand” in Bradshaw’s taxonomy.

3.7.2 Prospect Theory and Loss Aversion

Accepted definition. Kahneman and Tversky [1979], p. 263: “people underweight outcomes that are merely probable in comparison with outcomes that are obtained with certainty.” Losses loom larger than equivalent gains: loss aversion coefficient $\lambda \approx 2.25$.

Formal definition. Prospect theory value function: $v(x) = x^\alpha$ if $x \geq 0$; $-\lambda(-x)^\beta$ if $x < 0$; with $0 < \alpha, \beta < 1$ (diminishing sensitivity) and $\lambda > 1$ (loss aversion). Probability weighting: $w(\pi) = \pi^\gamma / (\pi^\gamma + (1 - \pi)^\gamma)^{1/\gamma}$, $\gamma \approx 0.65$, overweighting small and underweighting large probabilities. The *cumulative prospect theory* [Tversky and Kahneman, 1992] extends this to multi-outcome prospects and is consistent with first-order stochastic dominance.

Implications for health economics: patients overweight rare severe side effects; out-of-pocket costs feel more salient than equivalent insurance premium increases; framing a treatment as “10% mortality” vs. “90% survival” affects uptake (framing effect, McNeil *et al.* [1982]).

3.7.3 Present Bias and Time Inconsistency

Accepted definition. Strotz [1955], p. 165: an individual who is “myopic” discounts the near future more heavily than the far future, exhibiting *dynamic inconsistency*: plans formed today are abandoned tomorrow. Laibson [1997] formalised this via quasi-hyperbolic discounting.

Formal definition. Standard exponential: $U = u(c_0) + \delta u(c_1) + \delta^2 u(c_2) + \dots$ Quasi-hyperbolic:

$$U = u(c_0) + \beta \sum_{t=1}^T \delta^t u(c_t), \quad \beta < 1. \quad (21)$$

β is the present-bias parameter; δ is the long-run discount factor. Present bias causes over-consumption of unhealthy goods today (smoking, fast food, alcohol) despite strong stated preferences for future health, inconsistent with the standard Grossman model. *Nudges* [Thaler and Sunstein, 2008] are low-cost interventions that re-frame choices to exploit present bias constructively: automatic enrolment in pension savings, default healthy options in cafeterias.

3.7.4 Framing Effects and Reference Dependence

Accepted definition. Tversky and Kahneman [1981], p. 453: “the framing of an act or outcome may depend upon the verbal or social context in which it appears. Different frames for the same situation can yield different choices.”

Discussion. Framing effects violate the invariance axiom of expected utility. In clinical communication: “95% survival rate”

and “5% mortality rate” are logically equivalent but generate different patient choices [McNeil *et al.*, 1982]. In health policy: presenting a tax as a “bonus for healthy choices” versus a “penalty for unhealthy choices” affects compliance even at identical financial incentives. Reference dependence means that the status quo exerts disproportionate influence on choices (status quo bias, Samuelson and Zeckhauser [1988]), relevant to the design of opt-in versus opt-out organ donation systems.

3.8 Dynamic Economics

3.8.1 Intertemporal Choice and Social Discounting

Accepted definition. Fisher [1930], Ch. IV: saving and investment are choices between present and future consumption, with the interest rate as the price of time. Samuelson [1937]: the discounted utility model provides the standard framework.

Discussion. The social discount rate determines how much weight is given to future relative to present welfare. This is directly relevant to intergenerational justice: a high discount rate devalues future generations’ welfare. The Stern [2007] review used a near-zero discount rate for climate policy, implying strong obligations to future generations; Nordhaus [2008] used a market-implied rate, giving much less weight to the distant future. This disagreement is fundamentally a justice question, not a technical one.

Formal definition. Standard model: $U = \sum_{t=0}^T \beta^t u(c_t)$, $\beta = 1/(1+\rho)$. Euler equation: $u'(c_t) = \beta(1+r_{t+1})u'(c_{t+1})$. Ramsey optimal growth: $\rho = r - g\eta$, where r is the interest rate, g the growth rate of consumption, η the elasticity of marginal utility.

3.8.2 Dynamic Programming and Bellman Equation

Accepted definition. Bellman [1957], p. ix: “An optimal policy has the property that whatever the initial state and initial decision are, the remaining decisions must constitute an optimal policy with regard to the state resulting from the first decision.” This *principle of optimality* underlies all recursive methods in economics.

Formal definition. For a discrete-time problem with state $x \in \mathcal{X}$, control $u \in \mathcal{U}$, reward $r(x, u)$, transition $f(x, u)$, and discount β :

$$V(x) = \max_{u \in \mathcal{U}} [r(x, u) + \beta V(f(x, u))]. \quad (22)$$

The Bellman equation (22) is a functional equation for the value function $V : \mathcal{X} \rightarrow \mathbb{R}$. Under contraction mapping conditions, iteration of the Bellman operator $TV(x) = \max_u [r + \beta V(f)]$ converges to the unique fixed point V^* [Stokey and Lucas, 1989].

The POMDP (Partially Observable Markov Decision Process) extension [?]:

$$V(b) = \max_{a \in \mathcal{A}} \left[r(b, a) + \beta \sum_o P(o|b, a) V(\tau(b, a, o)) \right], \quad (23)$$

where b is the belief state (probability distribution over hidden states s), τ the belief update rule, and o observations. The Bellman equation for POMDPs is the formal foundation of the neuro-symbolic policy framework discussed in this paper’s broader research programme.

Connection to Pareto dynamics. In a social planning problem, the state x is the current allocation, the control u is the institutional policy (which symmetry-breaking direction to choose), and the value function $V^*(x)$ is the maximal achievable social welfare from state x . The Bellman equation expresses recursive welfare optimisation. Non-ergodic dynamics (Proposition 4.8) mean that $V^*(x)$ may depend on the initial state x_0 and the trajectory history, not just the current state x : the value function is path-dependent, and welfare analysis must be conducted on trajectory space, not state space.

4 Algebraic and Physical Interpretation

4.1 Symmetry Group of a Utility Function: Exact Construction

We now give a precise, example-grounded account of the symmetry structure of consumer preferences.

Modeling rationale. An indifference relation is an equivalence relation on X : it partitions the consumption space into classes of mutually indifferent bundles. Group theory is the mathematics of symmetry and equivalence: the orbits of a group action are exactly the equivalence classes. Modelling indifference curves as group orbits is therefore not an analogy but a direct identification — G_U consists exactly of those diffeomorphisms leaving U unchanged, whose orbits are exactly the indifference curves. This makes precise the intuition that moving along an indifference curve is a symmetry of preferences, and provides the entry point for all subsequent structure (connections, curvature, Noether charges).

Definition 4.1 (Symmetry group of U). The *symmetry group* of $U : X \rightarrow \mathbb{R}$ is the group of diffeomorphisms that preserve U :

$$G_U = \{\varphi \in \text{Diff}(X) : U \circ \varphi = U\}. \quad (24)$$

The connected component of the identity in G_U , denoted G_U^0 , is a Lie group. Its Lie algebra is $\mathfrak{g}_U = \{v \in \Gamma(TX) : \mathcal{L}_v U = v(U) = \nabla U \cdot v = 0\}$, the vector fields tangent to the level sets of U .

For $n = 2$ goods and a smooth utility U with $\nabla U \neq 0$ everywhere, $\mathfrak{g}_U$ is one-dimensional at each point, spanned by the vector $\xi(x) = J \nabla U(x)$, where $J = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}$ is the

standard symplectic matrix. The integral curves of ξ are exactly the level sets of U , giving $G_U^0 \cong \mathbb{R}$ locally.

Example 4.1 (Cobb–Douglas). Let $U(x_1, x_2) = x_1^\alpha x_2^\beta$, $\alpha, \beta > 0$. The vector field tangent to the level sets is $\xi = (\beta x_1) \partial_{x_1} + (-\alpha x_2) \partial_{x_2}$ (up to normalisation), since $\xi(U) = \beta x_1 \cdot \alpha x_1^{\alpha-1} x_2^\beta + (-\alpha x_2) \cdot \beta x_1^\alpha x_2^{\beta-1} = \alpha \beta U - \alpha \beta U = 0$. The one-parameter group generated by ξ is

$$\varphi_t(x_1, x_2) = (x_1 e^{\beta t}, x_2 e^{-\alpha t}), \quad t \in \mathbb{R}. \quad (25)$$

Verification: $U(\varphi_t(x)) = (x_1 e^{\beta t})^\alpha (x_2 e^{-\alpha t})^\beta = x_1^\alpha x_2^\beta e^{\alpha \beta t - \alpha \beta t} = U(x)$. ✓

The orbit of $x = (x_1, x_2)$ under G_U^0 is the indifference curve $\mathcal{I}_c = \{x : x_1^\alpha x_2^\beta = c\}$ through x . Since φ_t is injective and the Jacobian $|\partial \varphi_t / \partial x|$ is generically non-zero, the orbit is diffeomorphic to $\mathbb{R}$, consistent with $\dim G_U^0 - \dim \text{Stab}(x) = 1 - 0 = 1$ (Theorem 2.2).

Example 4.2 (CES). Let $U(x_1, x_2) = (\alpha x_1^\rho + \beta x_2^\rho)^{1/\rho}$. The tangent vector to the level sets is $\xi = (\beta x_2^\rho x_1) \partial_{x_1} - (\alpha x_1^\rho x_2) \partial_{x_2}$ (up to normalisation). For $\rho = 1$ (linear utility), the level sets are hyperplanes with constant normal (α, β) and $\xi = (\beta, -\alpha)$ is a constant vector field; the group action is translation.

Example 4.3 (Quasilinear). Let $U(x_1, x_2) = x_1 + v(x_2)$. The tangent vector is $\xi = \partial_{x_1}$, and the one-parameter group is pure translation in the x_1 direction: $\varphi_t(x_1, x_2) = (x_1 + t, x_2)$. This reflects the absence of income effects on good 2.

4.2 Utility as Conserved Hamiltonian

The identification of U with a Noether conserved charge is made rigorous via Hamiltonian mechanics.

Modeling rationale. We require dynamics on X that remain on level sets of U while generating the full indifference curve. Hamiltonian mechanics is the canonical theory of dynamics with a conserved quantity: the flow $X_H = J \nabla H$ preserves H identically because $\dot{H} = \nabla H \cdot J \nabla H = 0$ by skew-symmetry of J . Setting $H = U$ and choosing J as the standard symplectic matrix on $\mathbb{R}_+^2$ is the minimal canonical choice, requiring no structure beyond U itself. The Hamiltonian framework then delivers Noether’s theorem at no extra cost: U is automatically the conserved charge of time-translation symmetry, which is the formal geometric content of utility ordinality.

Proposition 4.1 (Orbit flow as Hamiltonian flow). *Let $X = \mathbb{R}_+^2$ with symplectic form $\omega = dx_1 \wedge dx_2$ (the standard area form). The Hamiltonian flow of $H = U$ with respect to ω is the orbit flow of G_U^0 on the level sets of U .*

Proof. Hamilton’s equations for $H = U$ are $\dot{x}_1 = \partial U / \partial x_2$, $\dot{x}_2 = -\partial U / \partial x_1$ (from $i_{X_H} \omega = dH$). The vector field $X_H = (\partial U / \partial x_2, -\partial U / \partial x_1) = J \nabla U$ is tangent to the level sets of U since $X_H(U) = \nabla U \cdot J \nabla U = 0$ (as J is skew-symmetric). This is exactly the generator ξ of G_U^0 . □

Corollary 4.1 (Utility as Noether Hamiltonian charge). *By Noether’s theorem in the Hamiltonian formulation, the time-translation symmetry of the Hamiltonian system $(X, \omega, H = U)$ yields $H = U$ itself as the conserved charge: $\dot{U} = \{U, U\}_\omega = 0$ along any trajectory. The ordinality of U corresponds to the freedom to relabel the value of the conserved charge by any strictly monotone function.*

Caveat 4.1. Proposition 4.1 requires $n = 2$. For $n > 2$ goods, the symplectic form on $\mathbb{R}^{2n}$ involves momenta p_i (Lagrange multipliers for budget constraints), and the Hamiltonian structure lives on the extended space T^*X . The orbit structure on X alone is then generated by the null space of the bordered Hessian (??). A complete $n > 2$ treatment requires the Dirac–Bergmann constrained Hamiltonian formalism [Dirac, 1964] and is left for future work.

4.3 Preference Fibre Bundle and Connection

Definition 4.2 (Preference fibre bundle). The *preference fibre bundle* is the smooth fibre bundle $(X, \mathbb{R}, \pi_U, G_U^0)$ where $\pi_U = U : X \rightarrow \mathbb{R}$ is the projection, each fibre $\pi_U^{-1}(c) = \mathcal{I}_c$ is an indifference curve, and the structure group is $G_U^0 \cong \mathbb{R}$ (locally).

Modeling rationale. Saying that X is partitioned into level sets of U captures the static orbit structure but not the relationships between level sets. The fibre bundle adds exactly the missing layer: a connection specifying how fibres at different utility levels relate to one another (how the preferred direction changes as utility changes), and a curvature measuring the failure of this relationship to be flat. These objects give a coordinate-free encoding of the MRS (ratio of connection components) and its variation (curvature), which are the economically central properties of preferences that a bare partition cannot capture. The bundle framework is the minimal structure that makes preference geometry tractable.

The vertical subbundle is $V_x = T_x \mathcal{I}_{U(x)} = \ker(dU(x)) = \{v \in T_x X : \nabla U(x) \cdot v = 0\}$.

The natural connection. Given a Riemannian metric g on X (for concreteness, the Euclidean metric $g_{ij} = \delta_{ij}$), the horizontal complement is $H_x = \text{span}(\nabla U(x))$, the gradient direction. The connection one-form is

$$\omega_U = \frac{dU}{\|\nabla U\|^2} = \sum_i \frac{\partial U / \partial x_i}{\|\nabla U\|^2} dx_i. \quad (26)$$

This is an $\mathbb{R}$ -valued one-form (the structure group $G_U^0 = \mathbb{R}$ is Abelian), and $\omega_U(\nabla U) = \|\nabla U\|^2 / \|\nabla U\|^2 = 1$ (correctly normalised).

Proposition 4.2 (MRS as ratio of connection components). *The marginal rate of substitution $\text{MRS}_{ij}(x)$ equals the ratio of the i -th and j -th components of the connection one-form ω_U :*

$$\text{MRS}_{ij}(x) = \frac{(\omega_U)_i}{(\omega_U)_j} = \frac{\partial U / \partial x_i}{\partial U / \partial x_j}. \quad (27)$$

Proof. Direct: $(\omega_U)_i = (\partial U / \partial x_i) / \|\nabla U\|^2$. The ratio eliminates the common denominator. $\square$

Economic interpretation. The connection one-form ω_U selects, at each point $x \in X$, the “most efficient direction” for increasing utility (the gradient). The MRS is the relative weight of each coordinate direction in this efficiency measure. The connection does not depend on the choice of monotone transformation of U since a transformation $\tilde{U} = f(U)$ replaces $\nabla \tilde{U} = f'(U)\nabla U$, and the ratio $(\partial \tilde{U} / \partial x_i) / (\partial \tilde{U} / \partial x_j) = (\partial U / \partial x_i) / (\partial U / \partial x_j)$ is unchanged. This is the geometric content of ordinality.

4.4 Bundle Curvature and Preference Non-Linearity

Modeling rationale. The question “how non-linear are preferences?” requires a coordinate-free, transformation-invariant answer. Bundle curvature provides it. For linear utility, ∇U is constant: the preferred direction never rotates as you move through X , so the connection is flat ($\Omega_U = 0$). For convex preferences, the MRS changes along each indifference curve: the preferred direction rotates. This rotation is exactly what $\Omega_U = d\omega_U$ measures — the failure of horizontal transport around an infinitesimal loop to close. Curvature is thus the invariant, transformation-invariant encoding of diminishing MRS, unchanged under monotone utility transformations and thus a genuinely ordinal concept.

Proposition 4.3 (Curvature of the preference connection). *The curvature of ω_U is the two-form $\Omega_U = d\omega_U$. In coordinates:*

$$(\Omega_U)_{ij} = \partial_i(\omega_U)_j - \partial_j(\omega_U)_i = \frac{\partial_i(\partial_j U / \|\nabla U\|^2) - \partial_j(\partial_i U / \|\nabla U\|^2)}{1} \quad (28)$$

which simplifies to

$$(\Omega_U)_{ij} = \frac{\partial^2 U}{\partial x_i \partial x_j} \cdot \frac{0}{\|\nabla U\|^4} + \frac{(\partial_j U)(\partial_k U)(\partial^2 U / \partial x_i \partial x_k) - (\partial_i U)(\partial_k U)(\partial^2 U / \partial x_j \partial x_k)}{\|\nabla U\|^4} \quad (29)$$

Proof. Compute $d\omega_U$ directly using the quotient rule. For a detailed derivation, see the appendix. $\square$

Key examples.

- **Linear utility** $U = \alpha_1 x_1 + \alpha_2 x_2$: $\nabla U = (\alpha_1, \alpha_2)$ is constant, so $\omega_U = \sum_i \frac{\alpha_i}{\|\alpha\|^2} dx_i$ is a closed form. $\Omega_U = d\omega_U = 0$.
- **Zero curvature:** perfect substitutes have a flat preference bundle.
- **Cobb–Douglas** $U = x_1^\alpha x_2^\beta$: $\partial U / \partial x_1 = \alpha x_1^{\alpha-1} x_2^\beta$ and $\partial U / \partial x_2 = \beta x_1^\alpha x_2^{\beta-1}$; these vary with x , so $d\omega_U \neq 0$.
- **Non-zero curvature:** convex preferences have a curved bundle.
- **Leontief** $U = \min(x_1, x_2)$: ∇U is not defined at $x_1 = x_2$. The connection is singular at the kink; curvature is concentrated on the diagonal as a distributional object.

Observation 4.1. The curvature Ω_U of the preference bundle measures the “non-flatness” of indifference curves: $\Omega_U = 0$ iff the level sets of U are hyperplanes, *i.e.* iff preferences are linearly separable (perfect substitutes). The diminishing MRS of convex preferences is geometrically the positive curvature of the indifference curves, encoded in $\Omega_U \neq 0$.

4.5 Budget Constraints as Symmetry Breaking

Modeling rationale. Before a budget constraint is imposed, the consumer is indifferent among all points on an indifference curve: the orbit has a continuous symmetry ($G_U^0 \cong \mathbb{R}$ acting freely). The budget constraint selects a unique point, reducing this symmetry to the trivial group. In quantum field theory, the same mechanism — a symmetric system selecting a preferred ground state — is called spontaneous symmetry breaking, and Goldstone’s theorem guarantees massless modes from the broken generators. The modelling choice is justified by this structural identity: the economic and physical mechanisms are formally identical, which is why the Goldstone theorem yields a genuine economic interpretation (price signals) rather than a forced analogy.

Proposition 4.4 (Budget constraint as symmetry breaking). *Let G_U^0 act on the indifference curve $\mathcal{I}_{c^*}$. A generic linear budget constraint $p \cdot x = w$ with $p \gg 0$ intersects $\mathcal{I}_{c^*}$ in exactly one point x^* under Assumption ?? (strict convexity). The stabiliser at x^* is $\text{Stab}(x^*) = \{e\}$ (trivial), so the residual symmetry after breaking is $H = \{e\} \subsetneq G_U^0 \cong \mathbb{R}$.*

Proof. Strict quasi-concavity of U implies each indifference curve is strictly convex. A hyperplane is tangent to a strictly convex curve in at most one point when the tangent direction (price ratio p_1/p_2) matches the slope of the curve (MRS) at that point, which by the strict convexity of preferences occurs uniquely on $\mathcal{I}_{c^*}$. The stabiliser is trivial because the only element of $G_U^0 \cong \mathbb{R}$ fixing x^* is the identity (the group acts freely). $\square$

Observation 4.2. The number of broken symmetry dimensions is $\dim G_U^0 - \dim H = 1 - 0 = 1$. By the Goldstone theorem, this predicts one massless mode. We now identify it with the price signal.

4.6 Price as Goldstone Covector

In field theory, Goldstone bosons are massless scalar

Modeling rationale. After the budget constraint selects x^* , two directions coexist in $T_{x^*} X$: the Goldstone direction $\xi(x^*)$ tangent to the indifference curve (zero marginal utility change), and the gradient direction (maximum utility increase). Prices $p = \lambda^{-1} \nabla U(x^*)$ point in the gradient direction: they measure the cost of departing from the indifference orbit. Goldstone modes in physics are the cost-free directions at a symmetry-broken ground state; prices are the cost-measuring covectors perpendicular to those directions. The identification is structural: prices emerge from the symmetry breaking imposed by

the budget constraint, encoding the exchange rates between goods imposed by scarcity.

particles associated with the broken generators of G . They live in the tangent space of the orbit G/H at the ground state. In the economic context, the broken generator is the vector field $\xi(x^*)$ tangent to $\mathcal{I}_{c^*}$ at x^* . The Goldstone “direction” is the direction along the indifference curve that the budget constraint prohibits.

The price vector p is a covector (an element of the dual space $(T_{x^*}X)^*$, not the tangent space), and is identified via the first-order condition (??):

$$p = \lambda^{-1} dU(x^*) = \lambda^{-1} \nabla U(x^*) \quad (\text{using Euclidean metric}), \quad (30)$$

where $\lambda > 0$ is the Lagrange multiplier (marginal utility of income).

Proposition 4.5 (Prices as Goldstone covectors). *The price covector $p = \lambda^{-1} dU(x^*)$ is the unique covector annihilating the Goldstone direction $\xi(x^*)$ up to a positive rescaling. Specifically, $p \cdot \xi(x^*) = \lambda^{-1} \nabla U(x^*) \cdot J \nabla U(x^*) = 0$ (since J is skew-symmetric).*

Proof. $\nabla U \cdot J \nabla U = \nabla U^T J \nabla U = 0$ since $J = -J^T$ implies $v^T J v = 0$ for all v . $\square$

Remark 4.1. Proposition 4.5 shows that prices are *dual* to the Goldstone direction in the precise sense of linear algebra: the price covector annihilates the tangent to the broken orbit. This is the correct statement of the “prices as Goldstone modes” analogy: prices are not vectors in the tangent space of the orbit but covectors in the cotangent space, related by the metric (or the Lagrange multiplier) to the gradient of utility. The relative price ratio $p_i/p_j = \text{MRS}_{ij}(x^*)$ encodes the broken symmetry direction.

Example 4.1 (Cobb–Douglas prices). For $U = x_1^\alpha x_2^\beta$, the Goldstone direction at $x^* = (x_1^*, x_2^*)$ is $\xi(x^*) = (\beta x_1^*, -\alpha x_2^*)$. The price covector $p = (p_1, p_2)$ satisfies $p_1 \beta x_1^* - p_2 \alpha x_2^* = 0$, i.e. $p_1/p_2 = \alpha x_2^*/(\beta x_1^*) = \text{MRS}_{12}(x^*)$, which is exactly the optimality condition (??). $\checkmark$

4.7 Market Equilibrium as a Dynamical Fixed Point

The Walrasian tatonnement is the gradient flow

Modeling rationale. A Walrasian equilibrium is a state where all markets clear simultaneously: no agent wants to change behaviour. This is definitionally a fixed point of any adjustment dynamics. The tatonnement is the natural dynamics: prices rise where demand exceeds supply and fall otherwise, which is gradient flow of the excess demand functional. Modelling equilibrium as a fixed point is therefore definitionally correct. The dynamical perspective adds value by addressing questions invisible to the static definition: stability (does the economy converge?), multiplicity (can multiple equilibria coexist?), and their absence (limit cycles, chaos [Saari, 1985]).

$$\dot{p}_i = \alpha_i z_i(p), \quad z_i(p) = \sum_j x_i^j(p, p \cdot \omega^j) - \omega_i^j, \quad (31)$$

where $\alpha_i > 0$. Equilibrium p^* satisfies $z(p^*) = 0$, i.e. p^* is a fixed point of (31).

Proposition 4.6. *Under gross substitutability ($\partial z_i / \partial p_j > 0$ for $i \neq j$), the tatonnement system has a unique globally asymptotically stable equilibrium p^* , with Lyapunov function $V(p) = \sum_i (p_i/p_i^* - 1 - \ln(p_i/p_i^*))$.*

Proof. Arrow et al. [1958]. $\square$

Without gross substitutability, equilibria may be multiple, non-unique, or the system may exhibit limit cycles or strange attractors [Saari, 1985]. The non-convergence of general equilibria is a fundamental result, not a pathology.

4.8 Social Path Integral in Euclidean Signature

Modeling rationale. The Euclidean rather than Minkowski path integral is chosen for two reasons. First, the Minkowski path integral $\int \mathcal{D}[x] e^{iS/\hbar}$ yields complex amplitudes, not probabilities: it does not directly model stochastic social dynamics. Second, the Euclidean version is rigorously justified by the Feynman–Kac formula as the transition density of a diffusion process with drift $-\nabla V$ and Gaussian noise $\sqrt{\hbar_{\text{soc}}}$, which is exactly the stochastic model of institutional dynamics. The parameter $\hbar_{\text{soc}}$ has a clear institutional interpretation: $\hbar_{\text{soc}} \rightarrow 0$ gives deterministic welfare optimisation; $\hbar_{\text{soc}} \rightarrow \infty$ gives pure Brownian drift with no institutional guidance.

Let $V : X \rightarrow \mathbb{R}$ be the *social disutility functional* ($V = -W$ where W is a social welfare function). The Euclidean social path integral gives the probability of the social state transitioning from x_i to x_f in “time” T under institutional noise $\hbar_{\text{soc}}$:

$$K(x_f, x_i; T) = \int_{x(0)=x_i}^{x(T)=x_f} \mathcal{D}[x] \exp\left(-\frac{1}{\hbar_{\text{soc}}} \int_0^T [\tfrac{1}{2} |\dot{x}|^2 + V(x)] d\tau\right). \quad (32)$$

By the Feynman–Kac formula, K is the transition density of the stochastic process

$$dx_\tau = -\nabla V(x_\tau) d\tau + \sqrt{\hbar_{\text{soc}}} dW_\tau, \quad (33)$$

where W_τ is a standard Brownian motion. This is a *Langevin equation*: the social state is driven toward the welfare optimum by drift $-\nabla V$ and perturbed by Gaussian noise of intensity $\sqrt{\hbar_{\text{soc}}}$.

The stationary distribution of (33) is the Boltzmann measure:

$$\mu^*(x) \propto \exp\left(-\frac{2V(x)}{\hbar_{\text{soc}}}\right), \quad (34)$$

which concentrates near the minima of V (maxima of W) as $\hbar_{\text{soc}} \rightarrow 0$, recovering the deterministic optimum. For $\hbar_{\text{soc}} > 0$, the system explores the welfare landscape, with higher- W states receiving exponentially higher probability mass.

Remark 4.2. The connection to simulated annealing [Kirkpatrick *et al.*, 1983] is direct: the acceptance probability $e^{-\Delta V/h_{\text{soc}}}$ in simulated annealing is the Boltzmann factor of (34). Decreasing h_{soc} over time (*cooling schedule*) concentrates the distribution on the global optimum.

4.9 Pareto Frontier as Boundary Set

Modeling rationale. The Pareto frontier is colloquially called a boundary, but the topological definition adds content. A topological boundary point cannot be determined from local information alone to lie inside or outside the dominated set: every neighbourhood contains both dominated and non-dominated allocations. This formalises the instability of Pareto-optimal states — any perturbation can push the system into the dominated interior — and connects Pareto optimality to the physics of critical and saddle points in the subsequent dynamical analysis.

Proposition 4.7 (Pareto frontier as topological boundary). *Let $D = \{x \in X : \exists y \text{ with } x \prec_P y\}$ be the Pareto-dominated set. Then $\mathcal{P} \subseteq \partial D$.*

Proof. Let $x^* \in \mathcal{P}$. By definition, $x^* \notin D$. Under local non-satiation (Assumption ??), for any open ball $B(x^*, \varepsilon)$ there exists $y \in B$ with $U_i(y) > U_i(x^*)$ for some i , but $U_j(y)$ may be smaller for $j \neq i$, so y may or may not dominate x^* . More precisely: for smooth utilities with generic initial endowments, the Pareto set has full measure in the boundary of D . Hence $x^* \in \overline{D} \cap \overline{D}^c = \partial D$. $\square$

4.10 Non-Ergodicity of Social Systems

Proposition 4.8 (Path-dependent social systems are non-ergodic). *If the stochastic process (33) is path-dependent*

Modeling rationale. *Equilibrium economics implicitly assumes ergodicity: without it, time averages need not converge to ensemble averages, and long-run equilibrium analysis loses its foundation. Path dependence — the documented fact that past institutional choices constrain future possibilities — directly violates the conditions for ergodicity: with path dependence, the limiting distribution depends on history, so the ergodic limit does not exist. Modelling social systems as non-ergodic is therefore not a theoretical preference but an empirical requirement imposed by the documented phenomena of institutional lock-in and historical contingency.*

(i.e. if the effective drift $-\nabla V$ depends on history $(x_\tau)_{\tau \leq t}$, then the process has no unique stationary distribution and is non-ergodic.

Proof. Ergodicity requires a unique invariant measure μ^* such that the empirical measure $\frac{1}{T} \int_0^T \delta_{x_t} dt \rightarrow \mu^*$ a.s. for all initial conditions. Path dependence means the law of x_T depends on the entire history, so the limit measure depends on the initial condition, contradicting uniqueness. $\square$

The conditions for Proposition 4.8 hold empirically for: technology adoption with network effects [Arthur, 1994], institutional evolution with complementarities [North, 1990], and financial markets with feedback between prices and expectations [?].

Corollary 4.2. *If social dynamics are non-ergodic, the Pareto frontier $\mathcal{P}$ is not generically reachable from all initial conditions, and welfare analysis based on the “potential Pareto improvement” standard (Kaldor–Hicks) is path-dependent.*

4.11 Duality Triangle: Roy, Shephard, and Hotelling as a Single Geometric Structure

The three duality relations.

Modeling rationale. Roy, Shephard, and Hotelling appear as three separate techniques in microeconomics texts. The Legendre–Fenchel transform reveals them as a single mathematical fact: the gradient of a convex value function recovers the optimal quantity vector. The reason this holds is that economic value functions (indirect utility, expenditure, profit) are convex or concave in prices by construction (as optima of linear objectives over feasible sets), and the Legendre transform is the operation that maps a convex function to its dual and recovers optima by differentiation. Understanding the Legendre unification explains not just why the three relations hold, but when analogous relations hold in any model where a value function is convex in a set of parameters.

Consumer theory and producer theory each possess a complete duality: an optimised value function, differentiated with respect to prices, recovers the optimal quantity choice. The three classical relations are:

$$\text{(Roy): } x_i^*(p, w) = -\frac{\partial v/\partial p_i}{\partial v/\partial w}, \quad (35)$$

$$\text{(Shephard consumer): } h_i(p, \bar{u}) = \frac{\partial e(p, \bar{u})}{\partial p_i}, \quad (36)$$

$$\text{(Hotelling): } y^*(p_y, w_z) = \frac{\partial \pi(p_y, w_z)}{\partial p_y}. \quad (37)$$

Algebraic interpretation: Legendre transform. The unifying structure is the *Legendre–Fenchel transform* [?]. For a convex function $f : \mathbb{R}^n \rightarrow \mathbb{R}$, the conjugate is $f^*(y) = \sup_x [\langle y, x \rangle - f(x)]$. The duality relations (35)–(37) are all instances of:

$$\nabla_p f^*(p) = x^*(p) \quad (\text{gradient of conjugate} = \text{primal optimum}), \quad (38)$$

where f is the primal objective and f^* the value function. Specifically:

- Consumer: $e(p, \bar{u})$ is the Legendre conjugate of the expenditure in p ; Shephard follows directly from (38).
- Producer: $\pi(p_y, w_z)$ is the conjugate of $-c$ in $(p_y, -w_z)$; Hotelling follows.

• Roy: $v(p, w)$ is related to e by $v(p, e(p, \bar{u})) = \bar{u}$; Roy is the implicit differentiation of this identity.

Proposition 4.9 (Duality triangle as fibre map). *The three value functions (v, e, π) are sections of the cotangent bundle $T^*\mathbb{R}_+^n$ in price space. Their gradients $(\nabla_p v, \nabla_p e, \nabla_{p_y} \pi)$ are the corresponding quantity demands and supplies, i.e. sections of the base manifold $\mathbb{R}_+^n$ (allocation space). Roy, Shephard, and Hotelling are three instances of the canonical projection $T^*\mathbb{R}_+^n \rightarrow \mathbb{R}_+^n$.*

Remark 4.3. The envelope theorem (Equation (15)) is the infinitesimal version of Proposition 4.9: the derivative of the value function with respect to a parameter equals the partial derivative of the Lagrangian at the optimum. All three duality relations are consequences of the envelope theorem applied to different primal–dual pairs.

4.12 Revealed Preference as Orbit Consistency

Algebraic interpretation. Revealed preference theory can be recast in terms of the orbit structure of the symmetry group G_U . A demand observation (p^k, x^k) says that x^k was chosen at prices p^k . WARP says: if x^k is revealed preferred to x^l (both were affordable at p^k), then x^l cannot be revealed preferred to x^k .

Modeling rationale. Revealed preference asks: when does observed choice behaviour come from a utility-maximising agent? The orbit language gives a geometric answer. Utility maximisation means the agent always moves to a weakly higher utility orbit when given the opportunity: choices visit orbits in a weakly monotone order. WARP prevents this orbit sequence from cycling: if x^k is chosen over available x^l (higher orbit than x^l), then x^l cannot be chosen over available x^k (would imply a lower orbit is higher, a cycle). GARP extends this acyclicity transitively. Afriat’s theorem then shows that GARP-consistent data always admit an explicit orbit structure (a utility function whose level sets are the relevant orbits): orbit consistency is necessary and sufficient for rational choice.

Proposition 4.10 (WARP as orbit non-crossing). *WARP is equivalent to the statement that the revealed preference binary relation R is acyclic on the set of observations $\{x^k\}$. In orbit terms: the sequence of chosen allocations $\{x^k\}$ must lie on a sequence of non-intersecting orbits of G_U with strictly increasing utility levels when the budget directions change monotonically.*

Proof sketch. Each $x^k \in \mathcal{I}_{c_k}$ for some c_k . WARP requires that if $p^k \cdot x^l \leq w^k$ (so x^l is in the same budget set as x^k) and $x^k \neq x^l$, then $p^l \cdot x^k > w^l$ (so x^k is not affordable at (p^l, w^l)). This prevents the orbit map $k \mapsto c_k$ from cycling. $\square$

Afriat’s theorem as orbit existence. Afriat [1967] shows that finite observations $\{(p^k, x^k)\}$ satisfying GARP are consistent with utility maximisation. The algebraic content is: GARP guarantees the existence of utility values $\{u_k\}$ and “shadow

prices” $\{\lambda_k > 0\}$ such that $u_j \leq u_k + \lambda_k p^k \cdot (x^j - x^k)$ for all k, j . These values define a piecewise-linear utility function whose level sets are the indifference curves consistent with the data. The group G_U in this case is piecewise-linear, with orbits being the level sets of the Afriat utility.

4.13 Risk Aversion as Utility Curvature and Stochastic Orbit Measures

Arrow-Pratt as second-order bundle curvature. For a scalar utility function $u : \mathbb{R}_+ \rightarrow \mathbb{R}$

Modeling rationale. Why does risk aversion correspond to curvature? Jensen’s inequality answers: $\mathbb{E}[u(x)] \leq u(\mathbb{E}[x])$ if and only if u is concave. Risk aversion means the agent prefers the certain expected value to a lottery with the same mean — exactly Jensen’s inequality for concave u . The Arrow-Pratt coefficient $r_A = -u''/u'$ quantifies this concavity normalised by slope: it is the curvature of the utility graph $(x, u(x))$ relative to its tangent direction, a dimensionless and monotone-transformation-invariant quantity. Modelling risk aversion as curvature is the direct geometric content of Jensen’s inequality, not a metaphor.

, concavity $u'' < 0$ is risk aversion. The Arrow-Pratt measure $r_A(x) = -u''(x)/u'(x)$ is the ratio of second to first derivatives, i.e. the normalised curvature of the graph of u .

Observation 4.3 (Risk aversion as curvature of 1-D bundle). Consider the trivial bundle $\pi : \mathbb{R} \times \mathbb{R}_+ \rightarrow \mathbb{R}_+$ with fibre coordinate $u(x)$. The connection one-form is $\omega = u' dx/u' = dx$ (trivially flat in the horizontal direction), but the fibre curvature $\kappa(x) = -u''(x)/u'(x) = r_A(x)$ measures how rapidly the utility gradient changes along the base. Risk neutrality ($u'' = 0$) corresponds to a flat utility graph; risk aversion ($u'' < 0$) to a concave (negatively curved) graph.

Expected utility as orbit measure. A lottery $L = (x_s, \pi_s)_s$ generates a probability measure $\mu_L = \sum_s \pi_s \delta_{x_s}$ on the consumption space X . The vNM expected utility $EU(L) = \int u d\mu_L$ is the integral of the utility function against the lottery measure. Under the orbit decomposition $X = \bigsqcup_c \mathcal{I}_c$, the lottery probability π_s is distributed across different orbits (utility levels $c_s = u(x_s)$). Risk aversion means that the agent prefers the “average orbit” $\mathcal{I}_{EU(L)}$ to the lottery over orbits.

Proposition 4.11 (Certainty equivalent as orbit projection). *The certainty equivalent $CE(L) = u^{-1}(EU(L))$ is the unique point in X lying on the orbit $\mathcal{I}_{EU(L)}$ that the agent is indifferent between and the lottery L . The risk premium $\rho = \mathbb{E}[x] - CE(L)$ measures the gap between the expected value orbit $\mathcal{I}_{u(\mathbb{E}[x])}$ and the certainty equivalent orbit $\mathcal{I}_{EU(L)}$, scaled by the inverse of the utility function.*

Prospect theory as asymmetric bundle geometry. The Kahneman-Tversky value function $v(x - r)$ is defined relative to a reference point r . The orbit structure now has two

distinct regions: $x > r$ (gains, concave v) and $x < r$ (losses, convex v). Loss aversion $\lambda \approx 2.25$ means the connection one-form ω_v has a discontinuity at $x = r$: the gradient $|v'|$ is approximately 2.25 times larger on the loss side. This is a curvature singularity (a distributional curvature) at the reference point, analogous to a conical singularity in a Riemannian manifold.

4.14 Intertemporal Choice as Dynamical System; Present Bias as Non-Autonomous Flow

Euler equation as fixed point. The Euler equation (12) can be written as:

Modeling rationale. Why model intertemporal choice as a dynamical system? The Euler equation is a difference equation in the consumption sequence: the optimality condition at each period links adjacent periods. A difference equation is by definition a discrete dynamical system; the steady state is by definition a fixed point. The dynamical perspective answers questions inaccessible to the static Euler-equation approach: is the steady state stable? Are there multiple steady states? For hyperbolic discounting, the discount factor changes at $t = 0$, creating a two-regime non-autonomous system whose time-inconsistency is a bifurcation: the optimal plan changes regime at $t = 0$, which is structurally visible only in the dynamical formulation.

$$\frac{u'(c_t)}{u'(c_{t+1})} = \beta(1 + r_{t+1}). \quad (39)$$

This is a fixed-point condition on the *consumption growth rate* $g_t = c_{t+1}/c_t$: in steady state, $u'(c)/u'(cg) = \beta(1 + r)$ defines g^* . For CRRA utility $u(c) = c^{1-\gamma}/(1-\gamma)$, the steady-state growth rate is $g^* = [\beta(1 + r)]^{1/\gamma}$, a fixed point of the map $g \mapsto [\beta(1 + r)/u'(cg)/u'(c)]^{1/\gamma}$.

The consumption path $\{c_t^*\}$ is the trajectory of the dynamical system

$$c_{t+1} = f(c_t; r_{t+1}, \beta), \quad (40)$$

with f given implicitly by (39). Competitive equilibrium corresponds to the stable manifold of this system near the steady state c^* .

Hyperbolic discounting as non-autonomous flow. Under quasi-hyperbolic discounting (21), the effective discount factor between t and $t + 1$ depends on the current period t :

$$\delta_{\text{eff}}(t, t + 1) = \begin{cases} \beta\delta & \text{if } t = 0 \\ \delta & \text{if } t \geq 1 \end{cases}. \quad (41)$$

This makes the dynamical system (40) *non-autonomous*: the map f changes at $t = 0$. The resulting time-inconsistency is a bifurcation in the dynamical sense: the optimal plan at $t = 0$ (using $\beta\delta$ for the next step) differs from the optimal plan at $t = 1$ (using δ for the next step, not $\beta\delta$). The agent's future

self constitutes a different “dynamical regime”, and the game between present and future selves is a multi-player fixed-point problem [Laibson, 1997].

Connection to institutional dynamics. Present bias in individual behaviour has a macroeconomic analogue: institutional short-termism. Governments, firms, and health systems systematically underinvest in prevention (long-run benefits) relative to treatment (short-run benefits), which is a collective present bias. In the social path integral (32), present bias enters through a time-dependent drift term $-\nabla V_t(x)$ where V_t weights near-future states more heavily than the socially optimal V .

4.15 Homothetic Preferences as Scale-Symmetric Bundles; Engel Curves as Radial Orbits

Modeling rationale. Homothetic preferences are, by definition, invariant under proportional scaling of all quantities: the shape of indifference curves is unchanged when the entire bundle is scaled by $\lambda > 0$. The group $\mathbb{R}_{>0}$ acting by $x \mapsto \lambda x$ is exactly the scale transformations of consumption space, so homotheticity means $G_U \supseteq \mathbb{R}_{>0}$: scale is a sub-symmetry of the preference symmetry group. This immediately implies that orbits (indifference curves) are mapped to orbits by scaling, so any ray through the origin meets each orbit once: the income-expansion path must be a ray. The scale-symmetry model is therefore not a convenient assumption but the precise algebraic content of the definition of homotheticity.

Proposition 4.12 (Homotheticity as scale symmetry). *A utility function U is homothetic if and only if the group of positive rescalings $\mathbb{R}_{>0}$ acts as a symmetry on the bundle structure: there exists a function $f : \mathbb{R}_{>0} \rightarrow \mathbb{R}_{>0}$ such that $U(\lambda x) = f(\lambda)U(x)$ for all $\lambda > 0$. Equivalently, $G_U \supseteq \mathbb{R}_{>0}$ (the symmetry group of U contains the scale group).*

Proof. Homotheticity means U is a monotone transformation of a linearly homogeneous function $\tilde{U}$: $U = h(\tilde{U})$, $\tilde{U}(\lambda x) = \lambda \tilde{U}(x)$. Then $U(\lambda x) = h(\lambda \tilde{U}(x)) = h(\lambda h^{-1}(U(x)))$, which defines $f(\lambda) \cdot U(x)$ after re-defining via h . The scale group $\lambda \mapsto \lambda x$ preserves the orbit structure. $\square$

Corollary 4.3 (Engel curves as radial orbits). *For homothetic preferences, the income-consumption curve (Engel curve at fixed prices) is a ray through the origin: $x^*(p, \lambda w) = \lambda x^*(p, w)$ for all $\lambda > 0$. This ray is an orbit of the scale group $\mathbb{R}_{>0}$ acting on X by scalar multiplication.*

The curvature of Engel curves therefore directly reflects whether preferences are homothetic (zero curvature, straight ray) or non-homothetic (non-zero curvature, bent ray). Non-homothetic preferences correspond to a non-trivial action of the income variable on the orbit structure.

4.16 The Grand Pareto Conditions as Triple Tangency of Manifolds

Modeling rationale. The three Pareto conditions $\text{MRS} = \text{MRTS} = \text{MRT}$ are conventionally derived separately as three

independent optimality requirements. The algebraic framework reveals their unity: all three say that a common covector (the price vector p^*) simultaneously annihilates the Goldstone direction of three different broken symmetries. This explains why the competitive price mechanism is informationally efficient: a single scalar signal enforces all three conditions simultaneously. Market failures occur precisely when no single price vector can simultaneously enforce all three — when the economy has more Goldstone modes than the price system can handle, as with externalities (unpriced third-party effects add extra Goldstone modes) or public goods (non-rivalry creates a Goldstone mode for each agent).

Geometric interpretation. The three Pareto conditions (17) express simultaneous tangency of three families of level sets at the social optimum:

1. Consumer indifference curves $\mathcal{I}_c^i$ (level sets of U^i) for all consumers i are mutually tangent: $\nabla U^i(x^i) = \mu^i \nabla U^j(x^j)$ (common MRS).
2. Isoquants Q_y^k (level sets of F^k) for all firms k are mutually tangent: common MRTS.
3. The production possibility frontier (level set of the aggregate transformation function $T(y_1, y_2) = 0$) is tangent to the common consumer indifference curve: $\text{MRS} = \text{MRT}$.

Proposition 4.13 (Grand Pareto as common cotangent). *At a Pareto optimum, there exists a price vector $p^* \in (T^*\mathbb{R}_+^n)$ such that: (i) $\nabla U^i(x^i) = \lambda^i p^*$ for all i (consumer optimality); (ii) $\nabla F^k(z^k) = \mu^k p^*$ for all k (firm optimality); (iii) $\nabla T(y) = -\nu p^*$ (PPF tangency). The common price vector p^* is the shared covector that simultaneously annihilates all the Goldstone directions of the broken indifference, isoquant, and transformation symmetries.*

Proof. At competitive equilibrium: consumers set $\nabla U^i = \lambda^i p^*$ (from (??)); firms set $\nabla F^k = \mu^k w^*$ and $\nabla T = -\nu p^*$ (from profit maximisation). The common price vector p^* is exactly the competitive equilibrium price, and the First Welfare Theorem (??) shows this coincides with the Pareto optimum. $\square$

Remark 4.4. This proposition shows that the Walrasian price mechanism is a *universal Goldstone covector*: a single element of cotangent space that simultaneously annihilates all broken symmetry directions across consumers, firms, and the production frontier. Market failure (externalities, public goods, monopoly) corresponds to situations where no single covector can simultaneously satisfy all three tangency conditions.

4.17 Short Run vs. Long Run as Constrained vs. Unconstrained Dynamics

Modeling rationale. The short-run versus long-run distinction is a dynamical statement: in the short run, some variables cannot adjust; in the long run, all can. Modelling this

as constrained versus unconstrained dynamics is natural because short-run dynamics are exactly the full gradient flow restricted to the submanifold of states reachable given fixed constraints. The envelope result $c^{SR} \geq c^{LR}$ follows immediately: a constrained optimum cannot outperform an unconstrained one. Institutionally: legal constraints, property rights, and social norms impose short-run dynamics; institutional reform removes the constraints and allows the system to explore the full allocation manifold. This makes the short-run / long-run distinction a statement about the dimensionality of the accessible submanifold, not just a horizon distinction.

Observation 4.4 (Short-run cost as constrained flow). The short-run cost $c^{SR}(w, y; \bar{z}_k)$ (with some inputs fixed at $\bar{z}_k$) is the solution to a constrained minimisation. It corresponds to the social dynamical system (33) with some state variables frozen: the drift $-\nabla V$ operates only in the subspace of adjustable variables. The long-run cost $c^{LR}(w, y)$ is the unconstrained minimum, corresponding to the full gradient flow.

The envelope result $c^{SR} \geq c^{LR}$ is a statement about constrained vs. unconstrained minima: any constrained optimum is at least as costly as the unconstrained. In dynamical terms: a system constrained to move in a lower-dimensional submanifold (the short-run subspace of adjustable inputs) cannot reach states that require motion in the full manifold.

In institutional economics, this generalises directly: institutions that fix certain social variables (legal constraints, property rights, constitutional rules) impose short-run constraints. The social system moves along the constrained dynamics; only institutional reform (“moving to the long run”) allows exploration of the full allocation space.

4.18 Nash Equilibrium as Fixed Point of Best-Response Correspondence

Fixed-point formulation. The best-response correspondence of player i is:

Modeling rationale. Nash equilibrium is defined as a state where no agent has a unilateral incentive to deviate: each is already playing a best response to others. This is word-for-word the definition of a fixed point of the joint best-response correspondence $\text{BR}(s^*) = s^*$. The fixed-point model is a direct formalisation of the concept, not an approximation. Kakutani’s theorem (rather than Brouwer’s) is required because best-response correspondences are set-valued when the agent is indifferent between multiple optimal responses; Kakutani handles upper hemicontinuous convex-valued correspondences, which is the correct regularity class for mixed strategies.

$$\text{BR}_i(s_{-i}) = \arg \max_{s_i \in S_i} u_i(s_i, s_{-i}). \quad (42)$$

A Nash equilibrium is a fixed point of the joint best-response correspondence $\text{BR} : \prod_i S_i \rightrightarrows \prod_i S_i$, $\text{BR}(s) = \prod_i \text{BR}_i(s_{-i})$.

Theorem 4.1 (Nash existence via Kakutani). *If each S_i is compact and convex and each u_i is continuous and quasi-concave in s_i , then BR satisfies the conditions of Kakutani's fixed-point theorem [Kakutani, 1941] and a Nash equilibrium exists.*

This parallels the Arrow-Debreu existence proof exactly: both use Kakutani's theorem, and both identify equilibrium as a fixed point of a best-response (or excess-demand) correspondence. The Walrasian tatonnement (31) is a special case where the strategy space is the price simplex and the best-response is the excess demand zero.

Prisoner's dilemma as orbit trapping. In the prisoner's dilemma, the Nash equilibrium (D, D) is a fixed point of the best-response dynamics but is Pareto-dominated by (C, C) . In orbit terms: (D, D) lies on a lower social welfare orbit than (C, C) . The system is “trapped” in the lower orbit because unilateral deviations (individual symmetry breaking) move each player to an even lower utility orbit. The prisoner's dilemma is thus a case where the game-theoretic symmetry breaking (each player optimising individually) selects a Pareto-inferior orbit — the opposite of the welfare-maximising symmetry breaking.

Repeated interaction allows escape from this trap: the Folk theorem (Equation (43) below) shows that the higher orbit (C, C) can be sustained when agents have sufficiently long time horizons.

4.19 Repeated Games as Discrete Dynamical Systems; Folk Theorem as Invariant Set

Dynamical systems interpretation. An infinitely repeated game is a discrete-time dynamical

Modeling rationale. A repeated game is structurally an infinite sequence of stage-game outcomes: a discrete-time trajectory in the action-profile space. The natural question for any dynamical system is the characterisation of its ω -limit set — the set of states the system can visit indefinitely. The Folk theorem characterises the ω -limit set of subgame perfect equilibrium trajectories in the repeated game. Modelling the repeated game as a dynamical system is not optional for understanding the Folk theorem: it is a statement about the invariant set of the dynamics, which has no static analogue. Trigger strategies are policies that make the cooperative orbit absorbing by threatening permanent exit to the Nash equilibrium orbit if any deviation occurs.

system on the space of histories $H = \prod_{t=0}^{\infty} A$ (sequences of action profiles). A strategy σ_i is a map from histories to actions. The equilibrium payoff set is:

$$V^* = \{v \in \mathbb{R}^m : v = (1-\delta) \sum_{t=0}^{\infty} \delta^t r(a_t), a_t \text{ generated by SPE}\} \quad (43)$$

The Folk theorem states that V^* converges to the set $\{v : v_i \geq \underline{v}_i, v \text{ feasible}\}$ as $\delta \rightarrow 1$, where $\underline{v}_i$ is the minimax payoff of player i .

Invariant set interpretation. In dynamical systems language, V^* is the ω -limit set of the equilibrium dynamics: the set of payoff vectors that can be sustained indefinitely. The Folk theorem says that as patience grows ($\delta \rightarrow 1$), this invariant set expands to include all individually rational feasible payoffs. Cooperation at the socially optimal orbit (C, C) is sustainable in the invariant set when δ is large enough.

Trigger strategies (“cooperate until defection, then punish forever”) are the simplest mechanism that sustains cooperation: defection triggers a jump to the Nash equilibrium orbit (D, D) permanently. The threat of orbit-downgrade deters defection.

4.20 Nash Bargaining as Geometric Mean Optimization on Orbit Space

Geometric interpretation. The Nash bargaining problem (18)

Modeling rationale. Why does the geometric mean emerge from Nash's axioms? Each axiom eliminates a degree of freedom in the class of possible solutions. Symmetry eliminates agent-specific weighting. Scale invariance eliminates dependence on the units of utility measurement (rescaling one agent's utility — relabelling their orbits — should not change the solution). Pareto optimality restricts the solution to the frontier. IIA restricts it to depend only on the shape of the feasible set near the solution. The geometric mean $\prod_i (u_i - d_i)$ is the unique functional on surplus vectors satisfying all four. Scale invariance is the key: it requires invariance under re-labelling of utility orbits, which uniquely selects the product form.

maximises $\prod_i (u_i - d_i)$ over the feasible set F . Setting $w_i = u_i - d_i$ (surplus relative to disagreement point), the objective $\prod_i w_i$ is the geometric mean of surpluses.

Observation 4.5 (Nash bargaining on utility orbits). The feasible set F is the image of the contract curve (Pareto set) in utility space. The disagreement point d corresponds to the orbit levels $\mathcal{I}_{d_i}$ for each agent. The Nash solution u^* is the point on the Pareto orbit that maximises the geometric mean of surpluses. On a log-scale, this is equivalent to maximising $\sum_i \ln(u_i - d_i)$, a concave functional over the convex utility possibility set.

The Nash bargaining solution has an axiomatic characterisation that parallels the symmetry structure of the orbit framework: the symmetry axiom requires invariance under re-labelling of agents (symmetry of the solution with respect to the permutation group); scale invariance requires invariance under affine utility transformations (which change the labels on orbits but not their ordering).

4.21 Edgeworth Box as Product Bundle; Contract Curve as Tangency Locus

Product bundle structure. In a two-agent exchange economy, the allocation space

Modeling rationale. The Edgeworth box is the Cartesian product of two agents' consumption spaces subject to a resource constraint. The product bundle is the natural mathematical model for two independent spaces coupled by a constraint, with each agent possessing their own preference bundle (connection, curvature, Goldstone structure). The contract curve as fibre tangency is then an immediate consequence: Pareto optimality requires the two bundles' fibre tangent spaces to coincide at the allocation point. This geometric derivation is more transparent than the classical MRS-equality argument because it exhibits the Pareto condition as a fibre intersection problem, making visible which geometric properties of each agent's preferences are involved.

is $X = X^A \times X^B$ subject to $x^A + x^B = \omega$. Each agent has their own preference bundle: $(X^A, \mathbb{R}, \pi_{U^A}, G_{U^A}^0)$ and $(X^B, \mathbb{R}, \pi_{U^B}, G_{U^B}^0)$.

Definition 4.3 (Exchange fibre bundle). The *exchange bundle* is the product bundle $\pi_A \times \pi_B : X \rightarrow \mathbb{R} \times \mathbb{R}$, with fibres $\mathcal{I}_{c_A}^A \times \mathcal{I}_{c_B}^B$ at each utility level pair (c_A, c_B) .

Proposition 4.14 (Contract curve as fibre intersection). The contract curve $\mathcal{C}$ is the set of points in X where the fibres of the two bundles are tangent:

$$\mathcal{C} = \{(x^A, x^B) \in X : T_{x^A} \mathcal{I}_{U^A(x^A)}^A = T_{x^B} \mathcal{I}_{U^B(x^B)}^B\} = \{(x^A, x^B) : \text{MRS}_A^A(x^A) = \text{MRS}_B^B(x^B)\} \quad (44)$$

Proof. Two indifference curves are tangent iff they have the same tangent space at the contact point, iff their gradient directions are collinear, iff their MRS values are equal. $\square$

A Walrasian equilibrium (p^*, x^{A*}, x^{B*}) is the unique point where the budget line (determined by p^*) intersects the contract curve. The First Welfare Theorem (??) says that this intersection is on $\mathcal{C}$, i.e. the equilibrium is where the two agents' Goldstone covectors $\lambda^A p^* = \nabla U^A(x^{A*})$ and $\lambda^B p^* = \nabla U^B(x^{B*})$ are both proportional to the same price covector p^* .

4.22 Arrow Impossibility as Categorical No-Go; SWF as Aggregation Functor

Categorical formulation. Let $\mathcal{P}$ be the category of preference orderings

Modeling rationale. The categorical language clarifies what Arrow's theorem is about. A social welfare function is not just any rule but must be systematic: it maps each profile of individual preference orderings to a social ordering in a way consistent with the structure of orderings (a functor). Arrow's conditions are different types of functorial conditions: Pareto is a preservation condition (the functor maps Pareto-improvements to

social improvements), IIA is a locality condition (the functor acts pointwise on pairs of alternatives), non-dictatorship is a non-degeneracy condition. Arrow's theorem says no functor of this type exists. This is a structural impossibility — not merely the non-existence of a convenient formula — which the categorical language makes precise.

(objects: rational preference relations on X ; morphisms: monotone transformations preserving the ordering). A social welfare functional $F : \mathcal{P}^m \rightarrow \mathcal{P}$ is a functor from the product category of individual preferences to a social preference.

Observation 4.6 (Arrow's theorem as functor non-existence). Arrow's impossibility theorem (Theorem 5.1) states that no functor $F : \mathcal{P}^m \rightarrow \mathcal{P}$ satisfies simultaneously: (i) Pareto (preserves unanimous preference); (ii) IIA (the social ordering between x and y depends only on the individual orderings of x vs. y); (iii) non-dictatorship. Conditions (i) and (iii) are naturality conditions on F ; condition (ii) is a locality condition (the functor acts pointwise on pairs). The impossibility reflects a fundamental obstruction in the category $\mathcal{P}$: there is no natural transformation from the product structure to a unified social ordering that is simultaneously local and non-dictatorial.

The Bergson–Samuelson SWF avoids Arrow's theorem by requiring cardinal utility and interpersonal comparisons, which amount to enriching the category $\mathcal{P}$ to a category $\mathcal{U}$ of cardinal utility functions with a specific metric structure. In this enriched category, the aggregation functor $W : \mathcal{U}^m \rightarrow \mathbb{R}$ exists (as a sum, minimum, or product), but at the cost of the additional structure.

4.23 Lorenz Curve as L^1 Distance; Stochastic Dominance as Partial Order on Distributions

Measure-theoretic interpretation. Let F be an income distribution

Modeling rationale. The L^1 interpretation of the Gini coefficient provides two insights unavailable from the standard formula. First, it gives the Gini a clear economic meaning: G is the normalised expected absolute pairwise income difference, stating directly what the Gini measures rather than just how it is computed. Second, it connects inequality measurement to the theory of L^p norms, enabling rigorous comparison of inequality indices as different norms on the distribution space. Stochastic dominance then provides the unambiguous comparison criterion: F is unambiguously less unequal than G iff the Lorenz curve of F lies above that of G , which is a partial-order statement on the space of distributions.

(CDF) and $L(p)$ the Lorenz curve. The Gini coefficient is:

$$G = 1 - 2 \int_0^1 L(p) \, dp = \frac{\iint |x - y| \, dF(x) \, dF(y)}{2\mu}, \quad (45)$$

where $\mu = \int x \, dF(x)$ is the mean. The right-hand side shows that G is the normalised expected absolute deviation between two randomly drawn incomes, i.e. the L^1 distance between two independent draws from F , normalised by mean income.

Observation 4.7 (Gini as L^1 metric on distribution space). The Gini coefficient G is a semi-metric on the space of income distributions: $G(F) \geq 0$, with $G = 0$ iff F is a point mass (perfect equality). It is not a proper metric (triangle inequality fails), but it is the L^1 earthmover distance between the Lorenz curve L and the diagonal of equality, normalised to $[0, 1]$.

Stochastic dominance as orbit partial order. First-order stochastic dominance (FSD) $F \succcurlyeq_1 G$ iff $F(x) \leq G(x)$ for all x . In orbit terms: FSD says that every percentile of F is at least as large as the corresponding percentile of G . FSD is a partial order on the space of distributions (reflexive, transitive, antisymmetric up to equality in distribution).

Second-order stochastic dominance (SSD) $F \succcurlyeq_2 G$ iff $\int_{-\infty}^x F(t) dt \leq \int_{-\infty}^x G(t) dt$ for all x (with equal means). SSD is equivalent to: $\mathbb{E}_F[u] \geq \mathbb{E}_G[u]$ for all concave u . In orbit terms: F SSD-dominates G iff every risk-averse agent prefers F to G , i.e. the Lorenz curve of F lies above that of G . This is a coarsening of the product partial order on the orbit space: SSD defines a social comparison that all concave (risk averse) utility functions agree on.

4.24 Optimal Taxation as Variational Problem on Orbit Space

Ramsey rule as constrained orbit selection. The Ramsey optimal taxation problem is

Modeling rationale. Taxes shift the effective price vector from p to $p + t$, changing the direction of the budget constraint and hence the symmetry-breaking direction selected by the consumer. The Ramsey problem asks which shift raises required revenue while minimising welfare loss — which direction of symmetry breaking is least costly. In orbit geometry, welfare loss is the drop in utility orbit caused by t ; revenue is the inner product $t \cdot h(p + t)$. Minimising the orbit displacement per unit revenue is the geometric content of the Ramsey problem. The inverse elasticity rule then has an intuitive geometric interpretation: inelastic Hicksian demand means steep indifference curves (small orbit displacement per unit price change), so taxing inelastic goods raises revenue with minimal orbit displacement.

: choose tax rates $\{t_i\}$ to maximise social welfare W subject to a revenue constraint $\sum_i t_i x_i^H(p + t, \bar{u}) \geq R$. The Lagrangian is:

$$\mathcal{L}_R = W - \nu \left[R - \sum_i t_i h_i(p + t) \right], \quad (46)$$

and the first-order conditions yield the Ramsey rule (20). In orbit terms: taxes shift the effective price vector from p to $p + t$, which changes the budget constraint hyperplane and thus the symmetry breaking direction. The Ramsey optimum is the tax vector that breaks symmetry in the direction that minimises the DWL orbit (the orbit of allocations with the least welfare loss relative to the pre-tax equilibrium).

Inverse elasticity rule as orbit curvature condition. The inverse elasticity rule $t_i/p_i \propto 1/|\varepsilon_{ii}^H|$ says: tax more heavily goods with inelastic Hicksian demand. In bundle geometry: inelastic demand corresponds to steep indifference curves (high MRS variation along the budget line), meaning the budget hyperplane crosses many utility orbits close together. Taxing these goods causes small orbit displacement per unit tax revenue. This is the geometric content of revenue efficiency.

4.25 Coase Theorem as Gauge Invariance of Efficient Outcomes

Observation 4.8 (Coase theorem as gauge invariance). The Coase theorem states: under zero transaction costs

Modeling rationale. The gauge-invariance analogy is justified by the structural parallel: in both gauge theory and the Coase setting, an internal degree of freedom (the gauge field / the property rights assignment) does not affect observables (physical forces / efficient allocation) when a condition holds (gauge symmetry / zero transaction costs). The property rights assignment is a gauge choice because it changes how initial endowments are allocated without changing productive possibilities. When transaction costs are zero, any reassignment is compensatable, making efficiency gauge-invariant. When costs are positive, the gauge fixing (property rights) affects efficiency: this is the analogue of a gauge-symmetry-breaking mass term. The Pigou–Coase distinction maps onto explicit versus spontaneous gauge fixing.

and perfect information, the efficient allocation is achieved regardless of the assignment of property rights. In gauge-theoretic language: the property rights assignment is a gauge choice (a choice of local section of the property rights bundle), and the efficient outcome is a gauge-invariant quantity that does not depend on this choice.

More precisely: let $\mathcal{G}$ be the group of permissible property rights assignments (“gauges”). For each gauge $g \in \mathcal{G}$, the equilibrium allocation $x^*(g)$ may differ, but if transaction costs are zero, $x^*(g) = x_{\text{eff}}^*$ for all g : efficiency is a gauge-invariant property.

When transaction costs are positive, the gauge dependence reappears: different property rights assignments yield different equilibria, and the “gauge fixing” (choice of property rights) affects the efficient outcome. This is the analogue of explicit symmetry breaking in gauge theory: the gauge invariance is approximate, broken by the transaction cost “mass term.”

Remark 4.5. The Pigou–Coase distinction maps onto the distinction between *explicit* symmetry breaking (Pigouvian taxes: the government imposes a specific gauge) and *spontaneous* symmetry breaking (Coasian bargaining: the agents select a gauge through their private agreement). Both achieve efficiency, but through different mechanisms of symmetry breaking.

4.26 Information Asymmetry as Obstruction to Orbit Separation

Adverse selection as type-space foliation. In the Akerlof market, sellers have types

Modeling rationale. The type-foliation language makes visible why information asymmetry causes market failure. Efficient markets require the price mechanism to resolve the type foliation: map each type θ to a distinct price so that buyers can infer quality. Under symmetric information, this resolution is possible (the map from types to prices is injective). Under asymmetric information, buyers observe only the projection of type space onto the price axis, which collapses the foliation: all types offering the same price pool together. Market failure is the topological collapse of this projection, and market unravelling is the extreme case where the projection collapses entirely — no price can separate types. This geometric perspective makes precise what “information resolves types” means and why its failure causes market failure.

$\theta \in [\underline{\theta}, \bar{\theta}]$. The full type space is a manifold Θ foliated by type levels. The *information asymmetry* means that the buyer can only observe the price p , not the type θ : the buyer’s information partition is coarser than the type foliation.

Observation 4.9 (Adverse selection as projection collapse). Market unravelling occurs when the projection from the full product space $X \times \Theta$ (allocation and type) to the observable space $X \times \{p\}$ (allocation and price) collapses the type foliation: at any given price p , only types $\theta \leq \theta^*(p)$ participate, and the buyer’s expected type falls as p falls, eventually making the market non-viable. This is a topological collapse: the pre-image $\pi^{-1}(p)$ in the type space shrinks to the empty set.

Moral hazard as constrained orbit selection. Under moral hazard, the agent selects effort e from a set E after the contract is signed. The principal cannot observe e ; only the outcome θ is observable. The second-best contract $(w(\theta), e^*)$ is the optimal point in the orbit of the contract space subject to the IC and IR constraints:

$$\text{IC constraint: } e^* \in \arg \max_{e'} \mathbb{E}[u(w(\cdot))|e'], \quad (47)$$

which restricts the feasible orbit of contracts. The principal-agent optimum is the point on the boundary of the feasible orbit that maximises the principal’s expected payoff.

Revelation principle as canonical orbit representative. The revelation principle (Myerson 1979) states that any equilibrium of any mechanism can be replicated by an incentive-compatible direct mechanism. In categorical terms: there is a functor from the category of mechanisms to the category of direct mechanisms, and this functor is essentially surjective (every IC outcome is achievable by a direct mechanism).

Direct mechanisms are the *canonical representatives* of equivalence classes of mechanisms under this functor: they are the “gauge-fixed” version of the mechanism design problem.

Signalling as orbit separation. In the Spence signalling model, the separating equilibrium is a fixed point of the signalling game where different types $\theta_H > \theta_L$ choose different signals $e_H > e_L$. The single-crossing condition ($\partial^2 c / \partial e \partial \theta < 0$) ensures that the orbit of type- θ_H agents in (e, w) space lies everywhere above the orbit of type- θ_L agents: the two type orbits are separated by the signal, allowing the market to identify types from signals.

4.27 Bounded Rationality and Behavioural Concepts as Perturbations of Optimal Orbits

Satisficing as orbit stopping rule. Simon’s satisficing agent stops searching when they find

Modeling rationale. Three behavioural departures are modelled using the orbit framework, each with a specific algebraic expression making precise what the departure is geometrically. *Satisficing*: Finding the highest achievable orbit is costly when search requires effort. The optimal stopping rule under search costs is to stop at the first orbit exceeding threshold u_0 : this is not irrational but optimal under bounded computational resources. The first passage time to orbit $\mathcal{I}_{u_0}$ is the formal expression of this stopping rule.

Loss aversion: The value function has a kink at r because hedonic valuation changes sign there: gains and losses are governed by different curvature regimes. The change of curvature sign at r is a singularity in the connection 1-form, correctly capturing the empirical discontinuity $|v'(r^-)| \approx 2.25|v'(r^+)|$.

Framing effects: Standard theory requires choice to depend only on the underlying lottery, not on its description: coordinate independence in outcome space. Framing effects violate this. Modelling this as coordinate non-invariance of ω_U formalises the violation: the connection transforms non-trivially under re-framing, as if the preferred direction changes with the coordinate system used to describe the outcomes.

an option with utility above an aspiration level u_0 :

$$\text{Stop at } x^* \text{ if } U(x^*) \geq u_0. \quad (48)$$

In orbit terms: the agent stops when they first enter the orbit $\mathcal{I}_{u_0}$ or higher. The aspiration level u_0 defines a target orbit, and satisficing is the dynamical process of searching until the orbit is reached. The optimal agent (standard neoclassical) searches until the highest achievable orbit is found; the satisficing agent stops at the first adequate orbit. The difference is a stopping time in the search dynamical system.

Loss aversion as asymmetric orbit geometry. In prospect theory, the reference point r creates an asymmetry in the utility orbit structure. Define two regions: the gain region $X^+ = \{x :$

$x > r\}$ and the loss region $X^- = \{x : x < r\}$. The value function $v(x - r)$ has different curvatures in the two regions: concave in X^+ (diminishing sensitivity to gains) and convex in X^- (diminishing sensitivity to losses), with a kink at $x = r$. The orbit structure has a singular point at $x = r$: the bundle curvature Ω_v has a delta-function singularity at the reference point, reflecting the discontinuous change in the gradient $|v'|$. Loss aversion means this singularity is asymmetric: $|v'(r^-)| \approx \lambda |v'(r^+)|$ with $\lambda \approx 2.25$.

Framing effects as non-invariant coordinate choice. In the optimal framework, the preference order $\succsim$ is invariant under re-labelling of outcomes: whether a health outcome is described as “10% mortality” or “90% survival” should not matter (both describe the same lottery). Framing effects [Tversky and Kahneman, 1981] violate this invariance: the choice changes with the description even when the underlying distributions are identical.

In geometric terms: framing is a choice of coordinates on the outcome space X . If the utility function U is frame-invariant, then $U(f(x)) = g(U(x))$ for any re-framing $f : X \rightarrow X$ (a monotone relabelling of outcomes). Framing effects mean U is not frame-invariant: the effective utility changes with the coordinate system. This is a violation of the gauge symmetry of the preference bundle: the connection one-form ω_U transforms non-trivially under re-framing, producing physical (choice) effects from a coordinate change.

4.28 Bellman Equation as Value Function Fixed Point; POMDP and Belief-State Geometry

Bellman operator as contraction. The Bellman operator $T : V \mapsto TV$ defined by

Modeling rationale. The Bellman equation is adopted because it is the necessary and sufficient condition for optimality in sequential decision problems. Bellman’s principle of optimality states that the value of being in state x equals the best immediate reward plus the discounted value of the optimal continuation: this is not an approximation but the definition of the value function. The contraction property ($\beta < 1$) makes V^* uniquely determined and computable by iteration. The connection to the social path integral is through the Legendre–Fenchel transform in the zero-noise limit: the path integral is dominated by the welfare-maximising trajectory, which the HJB equation (49) characterises as the value function. The two formulations are Legendre duals: one sums over all trajectories (path integral), the other characterises the optimal trajectory via a PDE (HJB).

$(TV)(x) = \max_u [r(x, u) + \beta V(f(x, u))]$ is a contraction on the space of bounded functions $(C_b(X), \|\cdot\|_\infty)$ with modulus $\beta < 1$ [Stokey and Lucas, 1989]. By the Banach fixed-point theorem, T has a unique fixed point $V^* = TV^*$, which is the Bellman equation (22). Value iteration $V_{n+1} = TV_n$ converges geometrically: $\|V_n - V^*\|_\infty \leq \beta^n \|V_0 - V^*\|_\infty$.

Observation 4.10 (Social Bellman equation). The social planning problem of Section 4.8 can be reformulated as a Bellman equation. The social state x evolves stochastically via the Langevin equation (33). The social value function $V^*(x) = \sup_{\text{policy}} \mathbb{E}[\int_0^\infty e^{-\rho t} W(x_t) dt | x_0 = x]$ satisfies the Hamilton–Jacobi–Bellman (HJB) equation:

$$\rho V^* = W(x) - \nabla V^* \cdot \nabla V + \frac{\hbar_{\text{soc}}}{2} \Delta V^*, \quad (49)$$

where Δ is the Laplacian. The HJB equation (49) is the continuous-time Bellman equation for the stochastic welfare problem, connecting the path integral formulation (Section 4.8) to the dynamic programming formulation.

POMDP and belief-state manifold. When the social state x is not fully observable, the planner maintains a belief state $b \in \Delta(\mathcal{X})$ (a probability distribution over possible states). The POMDP Bellman equation (23) operates on the *belief-state manifold* $\Delta(\mathcal{X})$, which has the structure of a statistical manifold equipped with the Fisher–Rao metric [Amari, 2016].

Observation 4.11 (Belief updates as geodesics on the statistical manifold). Bayesian belief updates $b \rightarrow \tau(b, a, o)$ (after action a and observation o) are flows on $\Delta(\mathcal{X})$. Under the Fisher–Rao metric, the Bayesian update is the unique update rule that is “closest to zero information gain” in the sense of minimising the KL divergence between the prior and posterior (Pythagoras’ theorem for information geometry [Amari, 2016]). The optimal POMDP policy is the fixed point of the Bellman operator on the belief-state manifold, generalising the social fixed-point analysis of Section 4.18 to the partial information setting.

Connection to generative justice. The POMDP framework provides a natural formalisation of the epistemic challenge in justice: a social planner who cannot fully observe the social state must form beliefs about it and take actions under uncertainty. The belief state b_t is the planner’s probability distribution over social states at time t ; generative justice (Definition 5.1) requires that the symmetry-breaking choice at each t be made in a way that maximises the expected justice score $\mathbb{E}_{b_t}[j(x_t, \dot{x}_t, G_t, H_t)]$ over the belief distribution.

4.29 Health State Space as Product Manifold; QALY and DALY as Measures

EQ-5D as product manifold. The EQ-5D health state space is the finite product

Modeling rationale. EQ-5D is defined by independent assessment of five dimensions: mobility, self-care, usual activities, pain/discomfort, anxiety/depression. Independence is the defining property of a product structure, so the health state space is indeed a product. The QALY weight $q : \mathcal{H} \rightarrow [0, 1]$ is a section of the trivial line bundle over $\mathcal{H}$, reducing the multi-dimensional space to a scalar suitable for decision analysis.

Different national tariffs correspond to different sections of the same bundle: the same total space $\mathcal{H}$ with different population preference weightings. The complementarity $\text{QALY} + \text{DALY} = t$ reflects a dual perspective on the same health trajectory: QALY integrates the quality weight q , DALY integrates the complement $1 - q$. They are dual sections of the health bundle, measuring the same trajectory from opposite orientations.

$\mathcal{H} = \{1, 2, 3, 4, 5\}^5$ (five dimensions, five levels each), a discrete approximation to a product manifold. The continuous analogue is:

$$\mathcal{H} = [0, 1]^5, \quad (50)$$

where each coordinate $h_k \in [0, 1]$ represents the k -th health dimension at full (1) to worst (0) function. The utility valuation function $q : \mathcal{H} \rightarrow [-0.285, 1]$ maps this product manifold to a scalar, reducing the multi-dimensional health state to a QALY weight.

Observation 4.12 (QALY weight as bundle section). The valuation function q is a section of the trivial bundle $\mathcal{H} \times \mathbb{R} \rightarrow \mathcal{H}$: it assigns a real number (utility weight) to each health state. Different value sets (UK tariff, US tariff, Dutch tariff) correspond to different sections of this bundle, encoding the preferences of different reference populations. The variation across tariffs is the “gauge freedom” in QALY valuation: different choices of reference population yield different sections.

QALY and DALY as different measures on orbit space.

The $\text{QALY} = q \times t$ and $\text{DALY} = \text{YLL} + \text{YLD} = N \times L + I \times DW \times L$ both integrate health quality weights over time, but use different orientations:

- *QALY*: measures health *gain*, integrating the weight q against a health intervention’s time effect. It is the signed area between the treatment and control health trajectories in (t, q) space.
- *DALY*: measures health *burden*, integrating the disability weight $DW = 1 - q$ against the duration of illness. It is the signed area between full health and the actual health trajectory.

In measure-theoretic terms: QALY and DALY are integrals of complementary measures on the health state orbit space. $\text{QALY} = \int_0^t q(s) \, ds$ and $\text{DALY} = \int_0^t (1 - q(s)) \, ds$, so $\text{QALY} + \text{DALY} = \text{total time } t$. Maximising QALYs is equivalent to minimising DALYs for fixed time horizon: the two metrics differ only in orientation.

Payment mechanisms as principal-agent contracts on provider orbit space. Each payment mechanism defines a mapping from provider actions (quantity, quality, efficiency) to provider revenue. Different mechanisms select different orbits in the provider welfare space:

$$\text{Revenue}(q, e) = P(q, e; \text{mechanism}), \quad (51)$$

where q is quantity and e is quality/effort. The principal (payer) designs P to align the provider’s optimal orbit (maximising provider welfare) with the social optimal orbit (maximising patient welfare and efficiency). Fee-for-service creates a revenue function that increases in q but is insensitive to e : the provider’s optimal orbit is high-quantity regardless of quality. Capitation creates a fixed revenue independent of q : the provider’s optimal orbit is cost-minimisation. Pay-for-performance ties revenue to quality indicators: the provider’s orbit is constrained to high- e regions.

The optimal payment mechanism is the solution to the principal’s maximisation problem subject to the provider’s incentive compatibility constraint, which is a mechanism design problem in the orbit space of provider welfare.

4.30 Updated Unified Correspondence Table

5 Ethics, Justice, and Formal Representations

5.1 Classical Ethical Frameworks

5.1.1 Virtue Ethics: Aristotle

Aristotle [350BCE] grounds ethics in *eudaimonia* (flourishing): the good life is excellent activity in accordance with virtue. Distributive justice (*dikaiaiosynê*) requires proportionality: $A/B = C/D$, where A and B are the goods received by two persons and C and D are their respective merits. Formally, an allocation x is distributively just if

$$\frac{x_i}{x_j} = \frac{m_i}{m_j} \quad \forall i, j, \quad (52)$$

where m_i is the “merit” of person i in the relevant dimension. The content of justice depends on the criterion of merit (need, contribution, status), which is itself contested.

5.1.2 Kantian Deontology

Kant [1785] grounds morality in reason. The categorical imperative in its universal law formulation: *act only on a maxim that you could will to be a universal law*. In the humanity formulation: *treat persons always as ends, never merely as means*.

The economic implication is a side-constraint on welfare maximisation: no allocation x^* is permissible if it treats any person merely as a means to aggregate welfare. Formally, the feasible set for the social planner is

$$X_K = \{x \in X : \text{no } i \text{ is treated merely as a means}\}, \quad (53)$$

and the planner maximises W over X_K . Specifying X_K in measurable terms is an open problem in applied ethics.

5.1.3 Utilitarianism

Bentham [1789] and Mill [1863] hold that the right action

maximises total welfare:

$$W_U(u_1, \dots, u_m) = \sum_{i=1}^m u_i \quad \text{or} \quad W_U = \sum_i \phi(u_i), \quad \phi'' < 0 \quad (54)$$

The concave version ($\phi'' < 0$) is *prioritarian*: it gives diminishing weight to increases in welfare at higher levels, favouring the worse off without the lexicographic priority of Rawls.

5.1.4 Rawlsian Contractarianism

Rawls [1971] derives principles of justice from the hypothetical agreement of rational agents behind a veil of ignorance. The two principles yield the *maximin* (more precisely *leximin*) welfare function:

$$W_R(u_1, \dots, u_m) = \min_i u_i, \quad (55)$$

which prioritises maximising the welfare of the worst-off agent. If two allocations have the same minimum, the second-lowest is compared, and so on (*lexicographic maximin*).

5.1.5 Libertarianism: Nozick

Nozick [1974] holds that distributions are just iff generated by just acquisitions and voluntary transfers, regardless of pattern. There is no simple welfare function representation; justice is a property of the process, not the outcome.

5.1.6 Capability Approach: Sen

Sen [1985, 1999] evaluates distributions over *capability sets*: $Q_i = \{f : f \text{ is achievable by person } i\}$. Justice requires: for each i , $Q_i \supseteq Q_{\min}$, where $Q_{\min}$ is a minimum threshold of capabilities (life, health, education, political participation). This cannot be reduced to utility maximisation because it evaluates freedom sets, not outcomes.

5.1.7 Luck Egalitarianism: Dworkin

Dworkin [2000] distinguishes brute luck (unchosen risk) from option luck (deliberate gamble). Justice requires neutralising brute luck. The *hypothetical insurance* mechanism: behind a veil of ignorance about their luck, all persons would purchase insurance against bad brute luck at actuarially fair rates. The just distribution is the post-insurance distribution that would result.

5.2 Justice in Economic Theory

5.2.1 The Fundamental Critique of Pareto Optimality

As established in ????, competitive markets generate Pareto-optimal allocations, but Pareto optimality is compatible with severe inequality. The central critique is:

Observation 5.1 (Pareto criterion as status quo bias). The set $\mathcal{P}$ of Pareto-optimal allocations depends on initial endowments (ω^i) . Different endowment distributions yield different $\mathcal{P}$. The Pareto criterion evaluates allocations conditional on endowments, without evaluating the endowments themselves. A society that begins from unjust endowments and achieves Pareto optimality through competitive markets has institutionalised, not corrected, the initial injustice.

5.2.2 Social Welfare Functions and Arrow's Theorem

Theorem 5.1 (Arrow Impossibility, 1951). *No social welfare function $F : \mathcal{R}^m \rightarrow \mathcal{R}$ (where $\mathcal{R}$ is the set of preference orderings) satisfies simultaneously: (i) Pareto efficiency; (ii) independence of irrelevant alternatives; (iii) non-dictatorship; (iv) unrestricted domain.*

Proof. Arrow [1951], Chapter V. $\square$

Arrow's theorem implies that any social welfare function must violate at least one normative condition. The choice of which condition to sacrifice is itself a justice decision: Rawls relaxes IIA by imposing cardinal utility comparisons; Nozick rejects the Pareto condition for non-entitlement transfers.

5.2.3 Justice Criteria and Welfare Functions

We summarise the major criteria:

Criterion	W or condition	Reference
Utilitarian	$\sum_i u_i$	Bentham [1789]
Prioritarian	$\sum_i \phi(u_i), \phi'' < 0$	Parfit [1984]
Rawlsian	$\min_i u_i$	Rawls [1971]
Capability	$Q_i \supseteq Q_{\min}$	Sen [1999]
Entitlement	Process condition	Nozick [1974]
Luck egalit.	Neutralise brute luck	Dworkin [2000]

5.3 Formal Representation of Justice in the Algebraic Framework

5.3.1 Terminal Justice as Fixed-Point Property

Utilitarian and Rawlsian justice both identify a target state $x^* \in X$ as the socially just allocation. In dynamical systems terms, x^* is a fixed point of the welfare gradient flow:

$$\dot{x} = \nabla W(x), \quad (56)$$

and justice is achieved when $x_t \rightarrow x^*$ as $t \rightarrow \infty$. The Lyapunov function for this flow is $V = -W$, and the flow converges if $-W$ is a Lyapunov function in a neighbourhood of x^* .

Critique. Proposition 4.8 shows that (56) need not converge when the drift depends on history. More fundamentally, the dynamical system (56) is a gradient flow only if W is a smooth function on X , which requires that individual utilities be cardinally comparable. This is precisely what Arrow's theorem problematises.

5.3.2 Procedural Justice as Gauge Symmetry

Nozick's historical entitlement theory can be formulated as a gauge-theoretic principle. Let $\mathcal{J}$ be the group of *just transformations*: those that respect all existing property rights and involve only voluntary exchange. A distribution x is just iff there exists a chain of just transformations from a just initial endowment ω^0 to x :

$$\text{Just}(x) \iff \exists \varphi_1, \dots, \varphi_k \in \mathcal{J} : x = \varphi_k \circ \dots \circ \varphi_1(\omega^0). \quad (57)$$

The set $\mathcal{O}_{\mathcal{J}}(\omega^0) = \mathcal{J} \cdot \omega^0$ is the orbit of the initial endowment under $\mathcal{J}$. Procedural justice holds that x is just iff $x \in \mathcal{O}_{\mathcal{J}}(\omega^0)$.

Formal limitation. The group $\mathcal{J}$ must be specified independently; it cannot be derived from within the framework. This is the gauge-freedom regress: the specification of which transformations are “just” is itself a justice question.

5.3.3 Capability Justice as Minimum Section

Sen's capability approach has a natural fibre-bundle representation. Let the fibre bundle $\pi : E \rightarrow X$ have fibres $E_x = Q_x$ (the capability set of the person at allocation x). A *section* of the bundle is a map $s : X \rightarrow E$ with $\pi \circ s = \text{id}$. The capability minimum condition requires:

$$s(x_i) \supseteq Q_{\min} \quad \forall i, \quad (58)$$

i.e. the section (capability profile) must lie above the minimum capability threshold $Q_{\min}$ at every agent position. The just allocation is any x for which such a section exists.

5.3.4 Generative Justice

Definition 5.1 (Generative justice). Let $\Gamma(X) = C([0, \infty), X)$ be the space of continuous social trajectories. A trajectory $\gamma \in \Gamma(X)$ satisfies *generative justice* with respect to a local scoring function j if

$$J(\gamma) := \int_0^\infty j(\gamma(t), \dot{\gamma}(t), G_t, H_t) dt > J_{\min}, \quad (59)$$

where G_t is the symmetry group at time t , H_t the residual symmetry after the budget-constraint selection at time t , and $J_{\min}$ is a normative threshold.

Remark 5.1. The integrand $j(\gamma(t), \dot{\gamma}(t), G_t, H_t)$ evaluates the normative quality of the symmetry-breaking event at time t : specifically, whether the budget-constraint selection H_t from the orbit $G_t \cdot \gamma(t)$ was made in a way that maximally expands the capability set of the least-advantaged agent. The integration over time means that the total justice score accumulates from all such events along the trajectory; no terminal state is privileged.

Proposition 5.1. *Generative justice $J(\gamma)$ is well-defined for any continuous trajectory, including those that do not converge*

to a fixed point (limit cycles, strange attractors, or divergent trajectories), provided j is bounded below and the integral converges.

Caveat 5.1. The construction of an explicit j that satisfies the normative desiderata (Rawlsian priority for the least-advantaged at each event, consistency with Arrow's theorem, measurability from observable data) is an open problem. We do not propose specific values for j in this paper.

Proposition 5.2 (Generative justice strictly generalises terminal justice). *Let W be a Rawlsian welfare function and $x^* = \arg \max_x W(x)$. If the social dynamical system converges to x^* (i.e. $\gamma(t) \rightarrow x^*$), then any trajectory converging to x^* that chose the least-advantaged-optimal breaking at every event has $J(\gamma) > J(\gamma')$ for any trajectory γ' that converges to x^* via less just breakings. In this sense, generative justice refines terminal justice on the set of convergent trajectories.*

Proof sketch. Two trajectories converging to the same x^* have the same terminal state, so terminal justice assigns them equal value. Generative justice distinguishes them by the integrals $J(\gamma)$ and $J(\gamma')$ of their local justice scores, which differ by assumption. Hence $J(\gamma) \neq J(\gamma')$ for trajectories that follow different paths to the same endpoint. $\square$

5.3.5 Connection to Biosemiotics and Universal Grammar

Kohn [2013], following Peirce [1931], argues that living systems are constitutively semiotic: sign processes are not uniquely human but pervade all living organisation. Forest mycorrhizal networks, immune systems, and economic markets all involve agents with incommensurable symbol systems exchanging information that they each interpret within their own representational framework.

In our algebraic language: different agents operate with different symmetry groups G_U^i (different utility functions, different preference structures). The *universal grammar* of the economic system would be a structure-preserving map (functor in the categorical sense [Lawvere, 1969]) between the preference categories of different agents, under which price signals are natural transformations: they map the optimal-breaking covector of one agent's preference bundle to that of another's in a compatible way.

Conjecture 5.1 (Universal economic grammar). There exists a category $\mathcal{E}$ of preference bundles and a functor $F : \mathcal{E} \rightarrow \mathcal{P}$ (price space) such that the price vector p at a Walrasian equilibrium is the unique natural transformation from all agents' optimal-breaking covectors to a common element of $\mathcal{P}$.

6 Discussion

6.1 Implications for Economic Theory

Price formation and market failure. If prices are Goldstone covectors of spontaneous symmetry breaking by budget

constraints, then administered prices (taxes, price controls, monopoly markups) correspond to *explicit* symmetry breaking that distorts the Goldstone direction. The welfare cost of such distortions is the DWL (??), which in this language is the curvature-squared integral over the deformed connection. A complete treatment would derive the DWL formula from the curvature of the deformed preference bundle.

Institutional rigidity and $\hbar_{\text{soc}}$. The Euclidean path integral (32) parametrises the trade-off between deterministic welfare optimisation and stochastic exploration. Low $\hbar_{\text{soc}}$ (rigid institutions) concentrates the system near the welfare optimum but may prevent escape from unjust local optima. High $\hbar_{\text{soc}}$ (weak institutions) allows exploration but dilutes the welfare signal. Calibrating $\hbar_{\text{soc}}$ from empirical data on policy variance is an open research question.

Ergodicity and the Kaldor–Hicks criterion. Peters [2019] demonstrates that the expected utility calculations underlying standard welfare economics systematically overestimate welfare gains when the growth process is non-ergodic. The geometric interpretation is that the Kaldor–Hicks compensation principle implicitly assumes that the Pareto frontier is reachable from any initial condition, which Proposition 4.8 shows is false for path-dependent systems. Policy evaluation should therefore be conducted on trajectory space, not state space.

6.2 Limitations and Open Problems

(1) *The group G_U for $n > 2$.* The explicit Hamiltonian structure of Proposition 4.1 requires $n = 2$. For $n > 2$, the orbit structure is $(n - 1)$ -dimensional but the symplectic form on $\mathbb{R}^{2n}$ requires momenta. The constrained Hamiltonian formalism [Dirac, 1964] provides the framework but has not been developed for the economic case.

(2) *Identification of the action functional.* The social Lagrangian $\mathcal{L}$ in (32) has not been derived from primitives. Relating $\mathcal{L}$ to observable choice data requires solving an inverse variational problem analogous to the identification problem in structural economics.

(3) *The local justice scoring function j .* Constructing j in Definition 5.1 that satisfies normative desiderata while remaining empirically estimable is the central open problem for the applied programme of generative justice.

(4) *Cardinal comparability.* The welfare functions of Section 5.1 require cardinal utility, which is inadmissible in the ordinal theory of ???. Resolving this requires either accepting interpersonal utility comparisons (using, e.g., von Neumann–Morgenstern utility under risk) or replacing utility with objectively measurable capabilities.

(5) *Soriette’s conjecture.* The conjecture in Section 5.3 that a universal economic grammar functor F exists has not been proved. A proof would require specifying the category $\mathcal{E}$ and constructing F explicitly.

7 Related Work

Physics-inspired economics. Mirowski [1989] documents the mechanical origins of neoclassical economics. Smith and Foley [2008] develops a path-integral approach using Minkowski signature (see our Section 4.8 for the Euclidean correction). Sornette [2014] surveys physical approaches to social systems. Mantegna and Stanley [1999] develops econophysics for financial markets.

Geometric approaches. Smale [1974] applied differential topology to general equilibrium. Chichilnisky [1985] proved topological social choice impossibility results. The fibre-bundle interpretation of preference structure is new to the present paper, though the connection between consumer theory and differential geometry has been noted informally by several authors.

Information geometry. Amari [2016] equips the space of probability distributions with a Riemannian structure via the Fisher metric. The connection between the preference bundle curvature Ω_U and the Fisher information metric on the demand correspondence deserves systematic investigation.

Dynamical systems in economics. Grandmont [1985] demonstrated chaotic dynamics in overlapping-generations models. Brock and Hommes [1997] and Farmer and Foley [2009] develop complex-systems approaches. Our contribution is the explicit connection between dynamical systems and the gauge-field geometry of preference and welfare structure.

Ergodicity and welfare. Peters [2019] provides the most rigorous recent argument for non-ergodicity in economic growth processes. Arthur [1994] and North [1990] document path dependence in technology and institutions.

Justice theory. Key texts are Rawls [1971], Nozick [1974], Sen [1999], Dworkin [2000]. Our generative justice proposal is most closely related to Derrida [1992] (justice as *a venir*) and Whitehead [1929] (process ontology). Teitel [2000] addresses justice in temporal transitions; generative justice extends this to all social dynamics.

Biosemitotics. Kohn [2013] and Peirce [1931] provide the semiotic foundations. The categorical universal-grammar conjecture is related to Lawvere [1969] and Fong and Spivak [2018].

8 Conclusion

We have developed a unified framework connecting consumer theory, group theory, gauge fields, Euclidean path integrals, and justice theory.

Established results. (I) The symmetry group G_U^0 of a utility function acts on the consumption space with indifference curves as orbits (Definition 4.1, verified for Cobb–Douglas, CES, quasilinear). (II) The utility function is the Hamiltonian of the orbit flow on (X, ω) and is a Noether conserved quantity under time-translation symmetry (Proposition 4.1). (III) Budget constraints break the orbit symmetry, selecting a unique optimum under strict convexity (Proposition 4.4). (IV)

The price vector is the Goldstone covector $\lambda^{-1} dU(x^*)$ that annihilates the broken Goldstone direction (Proposition 4.5). (V) The MRS is the ratio of components of the connection one-form ω_U on the preference fibre bundle (Proposition 4.2), and the bundle curvature vanishes for perfect substitutes and is generically non-zero for convex preferences (Proposition 4.3). (VI) The Euclidean social path integral is rigorously interpreted via the Feynman–Kac formula as the transition density of a Langevin stochastic process, with stationary Boltzmann measure concentrating near welfare optima (Section 4.8). (VII) Path-dependent social systems are non-ergodic; the Pareto frontier is not generically reachable from arbitrary initial conditions (Proposition 4.8).

Proposal. (VIII) Generative justice is defined as a functional $J(\gamma)$ on social trajectory space, evaluating the normative quality of each symmetry-breaking event (Definition 5.1). It strictly generalises terminal justice on convergent trajectories and remains well-defined for non-convergent ones.

Open problems. The key open problems are: the $n > 2$ Hamiltonian structure; the identification of the social Lagrangian; the construction of the local justice scoring function j ; and the proof or disproof of the universal economic grammar conjecture.

A broader aspiration of this work is to reground the physics-economics analogy on structurally sound foundations. Spontaneous symmetry breaking, Goldstone phenomena, and Euclidean path integrals are the appropriate tools not because economics resembles mechanics, but because both physics and economics confront systems whose aggregate behaviour is determined by symmetry rather than microscopic detail. Social and economic systems, which resist agent-level modelling but exhibit robust collective patterns, are natural candidates for the methods of modern theoretical physics, applied with the care and explicitness that the analogies demand.

Acknowledgements

These ideas emerged from a single conversation tracing a path from health economics to Pareto efficiency, from group orbits to Goldstone modes, from dynamical systems to ergodicity, and from symmetry breaking to generative justice. The observation that indifference curves are group orbits, and that MRS is a connection ratio, arose from a natural-science intuition about the geometry of level sets that proved more fruitful than anticipated.

References

- Akerlof, G. A. (1970). The market for lemons. *Quarterly Journal of Economics*, 84(3):488–500.
- Amari, S. (2016). *Information Geometry and Its Applications*. Springer, Tokyo.
- Aristotle (350 BCE). *Nicomachean Ethics*. Tr. T. Irwin (1999). Hackett, Indianapolis.

- Arrow, K. J. (1951). *Social Choice and Individual Values*. Wiley, New York.
- Arrow, K. J. and Debreu, G. (1954). Existence of an equilibrium for a competitive economy. *Econometrica*, 22(3):265–290.
- Arrow, K. J., Block, H. D., and Hurwicz, L. (1958). On the stability of the competitive equilibrium, II. *Econometrica*, 27(1):82–109.
- Arnold, V. I. (1989). *Mathematical Methods of Classical Mechanics*, 2nd ed. Springer, New York.
- Arthur, W. B. (1994). *Increasing Returns and Path Dependence*. University of Michigan Press.
- Baumol, W. J., Panzar, J. C., and Willig, R. D. (1982). *Contestable Markets*. Harcourt Brace Jovanovich.
- Bentham, J. (1789). *Introduction to the Principles of Morals and Legislation*. T. Payne, London.
- Brock, W. A. and Hommes, C. H. (1997). A rational route to randomness. *Econometrica*, 65(5):1059–1095.
- Chichilnisky, G. (1985). Von Neumann–Morgenstern utilities and cardinal preferences. *Math. Oper. Res.*, 10(4):633–641.
- Debreu, G. (1954). Representation of a preference ordering. In Thrall *et al.* (eds.), *Decision Processes*. Wiley, New York.
- Derrida, J. (1992). Force of law. In Cornell *et al.* (eds.), *Deconstruction and the Possibility of Justice*. Routledge, New York.
- Dirac, P. A. M. (1964). *Lectures on Quantum Mechanics*. Yeshiva University, New York.
- Drummond, M. F. *et al.* (2015). *Methods for the Economic Evaluation of Health Care Programmes*, 4th ed. Oxford University Press.
- Dummit, D. S. and Foote, R. M. (2004). *Abstract Algebra*, 3rd ed. Wiley.
- Dworkin, R. (2000). *Sovereign Virtue*. Harvard University Press.
- Evans, R. G. (1974). Supplier-induced demand. In Perlman (ed.), *Economics of Health and Medical Care*. Macmillan, London.
- Farmer, J. D. and Foley, D. (2009). The economy needs agent-based modelling. *Nature*, 460(7256):685–686.
- Farrell, M. J. (1957). The measurement of productive efficiency. *J. Roy. Stat. Soc. A*, 120(3):253–290.
- Feynman, R. P. and Hibbs, A. R. (1965). *Quantum Mechanics and Path Integrals*. McGraw-Hill.

- Fong, B. and Spivak, D. I. (2018). *An Invitation to Applied Category Theory*. Cambridge University Press.
- Goldstone, J. (1961). Field theories with superconductor solutions. *Nuovo Cimento*, 19(1):154–164.
- Grandmont, J. M. (1985). On endogenous competitive business cycles. *Econometrica*, 53(5):995–1045.
- Grossman, M. (1972). On the concept of health capital and the demand for health. *J. Polit. Econ.*, 80(2):223–255.
- Guckenheimer, J. and Holmes, P. (2013). *Nonlinear Oscillations, Dynamical Systems, and Bifurcations of Vector Fields*. Springer.
- Hall, B. C. (2015). *Lie Groups, Lie Algebras, and Representations*, 2nd ed. Springer.
- Holmstrom, B. (1979). Moral hazard and observability. *Bell J. Econ.*, 10(1):74–91.
- Kant, I. (1785). *Groundwork of the Metaphysics of Morals*. Tr. M. Gregor (1997). Cambridge University Press.
- Kirkpatrick, S., Gelatt, C. D., and Vecchi, M. P. (1983). Optimization by simulated annealing. *Science*, 220(4598):671–680.
- Kohn, E. (2013). *How Forests Think*. University of California Press.
- Lawvere, F. W. (1969). Adjointness in foundations. *Dialectica*, 23(3–4):281–296.
- Lee, J. M. (2013). *Introduction to Smooth Manifolds*, 2nd ed. Springer.
- Mantegna, R. N. and Stanley, H. E. (1999). *An Introduction to Econophysics*. Cambridge University Press.
- Mas-Colell, A., Whinston, M. D., and Green, J. R. (1995). *Microeconomic Theory*. Oxford University Press.
- McGuire, T. G. (2000). Physician agency. In Culyer and Newhouse (eds.), *Handbook of Health Economics*, vol. 1A. Elsevier, Amsterdam.
- Mill, J. S. (1863). *Utilitarianism*. Parker, London.
- Milnor, J. (1963). *Morse Theory*. Princeton University Press.
- Mirowski, P. (1989). *More Heat than Light*. Cambridge University Press.
- Nakahara, M. (2003). *Geometry, Topology and Physics*, 2nd ed. IOP Publishing.
- Nambu, Y. and Jona-Lasinio, G. (1960). Dynamical model of elementary particles. *Phys. Rev.*, 122(1):345–358.
- Nash, J. F. (1950). Equilibrium points in n -person games. *Proc. Nat. Acad. Sci.*, 36(1):48–49.
- Noether, E. (1918). Invariante Variationsprobleme. *Nachr. Ges. Wiss. Gött.*, pp. 235–257.
- North, D. C. (1990). *Institutions, Institutional Change and Economic Performance*. Cambridge University Press.
- Nozick, R. (1974). *Anarchy, State, and Utopia*. Basic Books.
- Pareto, V. (1906). *Manuale di Economia Politica*. Soc. Ed. Libreria, Milan.
- Parfit, D. (1984). *Reasons and Persons*. Oxford University Press.
- Peirce, C. S. (1931–1958). *Collected Papers*, 8 vols. Harvard University Press.
- Peskin, M. E. and Schroeder, D. V. (1995). *An Introduction to Quantum Field Theory*. Westview Press.
- Peters, O. (2019). The ergodicity problem in economics. *Nature Physics*, 15(12):1216–1221.
- Phelps, C. E. (2017). *Health Economics*, 5th ed. Routledge.
- Prigogine, I. and Stengers, I. (1977). *Self-Organization in Nonequilibrium Systems*. Wiley.
- Rawls, J. (1971). *A Theory of Justice*. Harvard University Press.
- Saari, D. G. (1985). Iterative price mechanisms. *Econometrica*, 53(5):1117–1131.
- Samuelson, P. A. (1938). A note on the pure theory of consumer's behaviour. *Economica*, 5(17):61–71.
- Samuelson, P. A. (1947). *Foundations of Economic Analysis*. Harvard University Press.
- Samuelson, P. A. (1954). The pure theory of public expenditure. *Rev. Econ. Stat.*, 36(4):387–389.
- Sen, A. (1985). *Commodities and Capabilities*. North-Holland.
- Sen, A. (1999). *Development as Freedom*. Oxford University Press.
- Simon, B. (1979). *Functional Integration and Quantum Physics*. Academic Press.
- Smale, S. (1974). Global analysis and economics, IIA. *J. Math. Econ.*, 1(1):1–14.
- Smith, E. and Foley, D. K. (2008). Classical thermodynamics and economic general equilibrium. *J. Econ. Dyn. Control*, 32(1):7–65.
- Sornette, D. (2014). Physics and financial economics. *Rep. Prog. Phys.*, 77(6):062001.

- Strogatz, S. H. (2018). *Nonlinear Dynamics and Chaos*, 2nd ed. Westview Press.
- Teitel, R. G. (2000). *Transitional Justice*. Oxford University Press.
- Tirole, J. (1988). *The Theory of Industrial Organization*. MIT Press.
- Varian, H. R. (1992). *Microeconomic Analysis*, 3rd ed. Norton.
- Walras, L. (1874). *Éléments d'Économie Politique Pure*. L. Corbaz, Lausanne.
- Walters, P. (1982). *An Introduction to Ergodic Theory*. Springer.
- Weinberg, S. (1995). *The Quantum Theory of Fields*, vol. 2. Cambridge University Press.
- Whitehead, A. N. (1929). *Process and Reality*. Macmillan.
- Willig, R. D. (1976). Consumer's surplus without apology. *Am. Econ. Rev.*, 66(4):589–597.
- Williams, A. (1997). Intergenerational equity. *Health Economics*, 6(2):117–132.
- Wilson, K. G. (1974). The renormalisation group. *Phys. Rep.*, 12(2):75–199.
- Yang, C. N. and Mills, R. L. (1954). Conservation of isotopic spin. *Phys. Rev.*, 96(1):191–195.
- Arrow, K. J. (1963). Uncertainty and the welfare economics of medical care. *Am. Econ. Rev.*, 53(5):941–973.
- Arrow, K. J. and Hahn, F. H. (1971). *General Competitive Analysis*. Holden-Day, San Francisco.
- Debreu, G. (1954). Valuation equilibrium and Pareto optimum. *Proc. Nat. Acad. Sci.*, 40(7):588–592.
- Edgeworth, F. Y. (1881). *Mathematical Psychics*. Kegan Paul, London.
- Fisher, I. (1927). *The Making of Index Numbers*. Houghton Mifflin, Boston.
- Hicks, J. R. (1939). *Value and Capital*, 2nd ed. Oxford University Press.
- Hicks, J. R. (1939). Foundations of welfare economics. *Econ. J.*, 49(196):696–712.
- Jevons, W. S. (1871). *The Theory of Political Economy*. Macmillan, London.
- Kahneman, D., Wakker, P. P., and Sarin, R. (2000). Back to Bentham? Explorations of experienced utility. *Q. J. Econ.*, 112(2):375–406.
- Kakutani, S. (1941). A generalisation of Brouwer's fixed point theorem. *Duke Math. J.*, 8(3):457–459.
- Kaldor, N. (1939). Welfare propositions of economics. *Econ. J.*, 49(195):549–552.
- Marshall, A. (1890). *Principles of Economics*. Macmillan, London.
- Mooney, G. H. (1987). What does equity in health mean? *World Health Stat. Q.*, 40(4):296–303.
- National Institute for Health and Care Excellence (2013). *Guide to the Methods of Technology Appraisal*. NICE, London.
- Osborne, M. J. and Rubinstein, A. (1994). *A Course in Game Theory*. MIT Press.
- Pigou, A. C. (1920). *The Economics of Welfare*. Macmillan, London.
- Rothschild, M. and Stiglitz, J. E. (1976). Equilibrium in competitive insurance markets. *Q. J. Econ.*, 90(4):629–649.
- Samuelson, P. A. (1954). The pure theory of public expenditure. *Rev. Econ. Stat.*, 36(4):387–389.
- Scitovsky, T. (1941). A note on welfare propositions in economics. *Rev. Econ. Stud.*, 9(1):77–88.
- Sen, A. K. (1970). *Collective Choice and Social Welfare*. Holden-Day, San Francisco.
- Slutsky, E. (1915). Sulla teoria del bilancio del consumatore. *Giornale degli Economisti*, 51:1–26.
- Smith, A. (1776). *An Inquiry into the Nature and Causes of the Wealth of Nations*. W. Strahan and T. Cadell, London.
- von Neumann, J. and Morgenstern, O. (1944). *Theory of Games and Economic Behavior*. Princeton University Press.
- Afriat, S. N. (1967). The construction of utility functions from expenditure data. *Int. Econ. Rev.*, 8(1):67–77.
- Arrow, K. J. (1965). *Aspects of the Theory of Risk-Bearing*. Yrjö Jahnssonin Säätiö, Helsinki.
- Atkinson, A. B. (1970). On the measurement of inequality. *J. Econ. Theory*, 2(3):244–263.
- Bellman, R. (1957). *Dynamic Programming*. Princeton University Press.
- Bergson, A. (1938). A reformulation of certain aspects of welfare economics. *Q. J. Econ.*, 52(2):310–334.
- Bradshaw, J. (1972). A taxonomy of social need. *Problems and Progress in Medical Care*, 7:69–82.

- Charnes, A., Cooper, W. W., and Rhodes, E. (1978). Measuring the efficiency of decision making units. *Eur. J. Oper. Res.*, 2(6):429–444.
- Coase, R. H. (1960). The problem of social cost. *J. Law Econ.*, 3:1–44.
- Condorcet, M. J. A. N. (1785). *Essai sur l'Application de l'Analyse à la Probabilité des Décisions Rendues à la Pluralité des Voix*. Imprimerie Royale, Paris.
- Devlin, N. J., Shah, K. K., Feng, Y., Mulhern, B., and van Hout, B. (2018). Valuing health-related quality of life: An EQ-5D-5L value set for England. *Health Economics*, 27(1):7–22.
- Ellis, R. P. and McGuire, T. G. (1986). Provider behavior under prospective reimbursement. *J. Health Econ.*, 5(2):129–151.
- Engel, E. (1857). Die Productions- und Consumtionsverhältnisse des Königreichs Sachsen. *Zeitschrift des Statistischen Bureaus des Königlich Sächsischen Ministeriums des Innern*, 8–9.
- EuroQol Group (1990). EuroQol—a new facility for the measurement of health-related quality of life. *Health Policy*, 16(3):199–208.
- Fisher, I. (1930). *The Theory of Interest*. Macmillan, New York.
- Friedman, J. W. (1971). A non-cooperative equilibrium for supergames. *Rev. Econ. Stud.*, 38(1):1–12.
- Gini, C. (1912). *Variabilità e Mutabilità*. Tipografia di Paolo Cuppini, Bologna.
- Hotelling, H. (1932). Edgeworth's taxation paradox. *J. Polit. Econ.*, 40(5):577–616.
- Houthakker, H. S. (1950). Revealed preference and the utility function. *Economica*, 17(66):159–174.
- Jonsen, A. R. (1986). Bentham in a box. *Law Med. Health Care*, 14(3–4):172–174.
- Kahneman, D. and Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–292.
- Laibson, D. (1997). Golden eggs and hyperbolic discounting. *Q. J. Econ.*, 112(2):443–478.
- Lerner, A. P. (1934). The concept of monopoly and the measurement of monopoly power. *Rev. Econ. Stud.*, 1(3):157–175.
- Lorenz, M. O. (1905). Methods of measuring the concentration of wealth. *Publ. Am. Stat. Assoc.*, 9(70):209–219.
- Luce, R. D. and Raiffa, H. (1957). *Games and Decisions*. Wiley, New York.
- Manning, W. G. *et al.* (1987). Health insurance and the demand for medical care. *Am. Econ. Rev.*, 77(3):251–277.
- McNeil, B. J., Pauker, S. G., Sox, H. C., and Tversky, A. (1982). On the elicitation of preferences for alternative therapies. *New Engl. J. Med.*, 306(21):1259–1262.
- Milgrom, P. and Segal, I. (2002). Envelope theorems for arbitrary choice sets. *Econometrica*, 70(2):583–601.
- Mirrlees, J. A. (1971). An exploration in the theory of optimum income taxation. *Rev. Econ. Stud.*, 38(2):175–208.
- Murray, C. J. L. and Lopez, A. D. (1994). Quantifying disability: Data, methods and results. *Bull. World Health Organ.*, 72(3):429–445.
- Myerson, R. B. (1979). Incentive compatibility and the bargaining problem. *Econometrica*, 47(1):61–73.
- Nash, J. F. (1950). The bargaining problem. *Econometrica*, 18(2):155–162.
- Nordhaus, W. D. (2008). *A Question of Balance*. Yale University Press.
- Pratt, J. W. (1964). Risk aversion in the small and in the large. *Econometrica*, 32(1–2):122–136.
- Ramsey, F. P. (1927). A contribution to the theory of taxation. *Econ. J.*, 37(145):47–61.
- Roy, R. (1947). La distribution du revenu entre les divers biens. *Econometrica*, 15(3):205–225.
- Samuelson, P. A. (1937). A note on measurement of utility. *Rev. Econ. Stud.*, 4(2):155–161.
- Samuelson, W. and Zeckhauser, R. (1988). Status quo bias in decision making. *J. Risk Uncertainty*, 1(1):7–59.
- Selten, R. (1965). Spieltheoretische Behandlung eines Oligopolmodells mit Nachfrageträgheit. *Zeitschrift für die gesamte Staatswissenschaft*, 121:301–324.
- Simon, H. A. (1955). A behavioral model of rational choice. *Q. J. Econ.*, 69(1):99–118.
- Spence, M. (1973). Job market signaling. *Q. J. Econ.*, 87(3):355–374.
- Stern, N. (2007). *The Economics of Climate Change: The Stern Review*. Cambridge University Press.
- Stinnett, A. A. and Mullahy, J. (1998). Net health benefits: A new framework for the analysis of uncertainty in cost-effectiveness analysis. *Med. Decis. Making*, 18(2 Suppl):S68–S80.
- Stokey, N. L. and Lucas, R. E. (1989). *Recursive Methods in Economic Dynamics*. Harvard University Press.

- 3016 Strotz, R. H. (1955). Myopia and inconsistency in dynamic
3017 utility maximization. *Rev. Econ. Stud.*, 23(3):165–180.
- 3018 Thaler, R. H. and Sunstein, C. R. (2008). *Nudge*. Yale Univer-
3019 sity Press.
- 3020 Tucker, A. W. (1950). A two-person dilemma. Unpublished
3021 notes; reprinted in Straffin, P. D. (1980), *The Prisoner's*
3022 *Dilemma*. *UMAP J.*, 1:101–103.
- 3023 Tversky, A. and Kahneman, D. (1981). The framing of deci-
3024 sions and the psychology of choice. *Science*, 211(4481):453–
3025 458.
- 3026 Tversky, A. and Kahneman, D. (1986). Rational choice and
3027 the framing of decisions. *J. Bus.*, 59(4):S251–S278.
- 3028 Tversky, A. and Kahneman, D. (1992). Advances in prospect
3029 theory. *J. Risk Uncertainty*, 5(4):297–323.
- 3030 von Neumann, J. (1928). Zur Theorie der Gesellschaftsspiele.
3031 *Math. Ann.*, 100(1):295–320.
- 3032 Working, H. (1943). Statistical laws of family expenditure. *J.*
3033 *Am. Stat. Assoc.*, 38(221):43–56.

Table 1: Extended correspondence between field theory, algebraic structures, and economic theory.

Field / Algebra	Structure	Economic concept
Symmetry group G_U^0	Lie group action on X	Indifference symmetry
G -orbit	Level set of U	Indifference curve
Noether charge	Hamiltonian $H = U$	Utility level
Scale group $\mathbb{R}_{>0}$	Scale invariant bundle	Homothetic preferences
Radial orbit	Ray through origin	Engel curve
SSB: $G \rightarrow \{e\}$	Budget constraint breaks orbit	Consumer optimum
Goldstone covector	$dU(x^*)$	Price vector p
Triple tangency	Common covector p^*	Grand Pareto conditions
Curvature Ω_U	Non-flat bundle	Diminishing MRS
Singular curvature at r	Reference-point kink	Loss aversion / prospect theory
Legendre transform	Value function gradient	Roy / Shephard / Hotelling
Envelope theorem	Chain rule at optimum	All duality relations
Orbit consistency (GARP)	Acyclic orbit visits	Revealed preference / Afriat
Fixed point of BR	Kakutani theorem	Nash equilibrium
Orbit trapping	Sub-optimal fixed point	Prisoner's dilemma
Invariant set as $\delta \rightarrow 1$	ω -limit set	Folk theorem / cooperation
Geometric mean on orbit	Nash product	Nash bargaining
Product bundle $\pi_A \times \pi_B$	Two-agent fibres	Edgeworth box
Fibre tangency locus	$MRS^A = MRS^B$	Contract curve
Common Goldstone covector	Universal p^*	Walrasian equilibrium
Categorical no-go	Functor obstruction	Arrow impossibility
L^1 norm on $[0, 1]$	Distance from diagonal	Gini coefficient
Partial order on CDFs	Stochastic dominance	Social welfare ranking
Gauge invariance	Outcome independent of gauge	Coase theorem
Explicit gauge breaking	Tax t shifts price	Pigouvian correction
Gauge choice	Property rights assignment	Initial endowment
Projection collapse	Type foliation collapse	Adverse selection
IC constraint on orbits	Feasible contract orbit	Moral hazard
Canonical functor	Direct mechanism rep.	Revelation principle
Separated orbits	Single-crossing condition	Signalling equilibrium
Gradient approximation	Stopping at aspiration orbit	Bounded rationality
Asymmetric bundle kink	$ v'(r^-) \approx \lambda v'(r^+) $	Loss aversion
Non-autonomous flow	Time-varying discount rate	Hyperbolic discounting
Coordinate non-invariance	Frame-dependent ω_U	Framing effects
Bellman contraction	Fixed point of T	Dynamic programming
HJB equation	Stochastic Bellman	Social path integral
Belief-state manifold	Statistical manifold $\Delta(X)$	POMDP / partial information
Geodesic belief update	Pythagoras information	Bayesian updating
Product manifold $[0, 1]^5$	Health state space	EQ-5D
Bundle section	Valuation tariff	QALY weight function
Complementary measures	QALY + DALY = time	Health evaluation duality
Orbit-revenue contract	Principal-agent on orbits	Payment mechanism design