A TREATISE

RENEWAL LIBERTARIANISM

substrate decay,

cyclical capture, &
the structural limits of liberty

PART ONE

TECHNICAL FOUNDATIONS

ANDREW A. GRATHWOHL

May 4, 2026

Contents

Introduction

The Operational Problem of Voluntary Association

Voluntary association is the ideal state for interactions among persons. This is not a narrowly libertarian commitment. It is the foundational premise shared by anarcho-capitalism, minarchism, classical liberalism, and the broader tradition of thought that takes voluntary cooperation between consenting persons as the standard against which interactions ought to be measured. Where these traditions differ is not in the foundational commitment but in their account of what the commitment requires institutionally. Anarcho-capitalism holds that voluntary association is best protected by the absence of any coercive substrate, since substrate by its nature interferes with voluntary association more than it enables it. Minarchism holds that voluntary association requires a minimal substrate confined to protective functions that voluntary actors cannot reliably perform for themselves. The framework presented here shares the foundational commitment with both alternatives and parts company with them on a different question: not whether substrate is necessary, but what happens to substrate over operational time once it exists and begins to function.

The question matters because the answer is not what either alternative tradition supposes. Anarcho-capitalism's answer is that the question does not arise, because substrate is unnecessary and should not exist. Module 1 of this framework demonstrates that this answer is wrong: voluntary cooperation in interactions whose stakes exceed the discounted future value of cooperative reputation cannot be sustained without external enforcement, and at the scales of modern political life, most interactions exceed this threshold. Without substrate, three failure modes emerge: defection cascade, in which cooperation in high-stakes interactions collapses; consolidation, in which a substrate-equivalent emerges through private accumulation of enforcement capacity; and substrate smuggling, in which informal substrate-like arrangements develop that evade the analytical attention of theory but perform substrate's functions in practice. None of these is consistent with the voluntary-association premise the anarcho-capitalist position is supposed to defend.

Minarchism's answer is that substrate, once established at its protective minimum through a properly designed founding moment, will remain at that minimum so long as constitutional vigilance is maintained. The framework's analytical machinery, developed across Modules 2 through 7, demonstrates that this answer is also wrong. Substrate is not a passive enforcer that occupies the role assigned to it and remains there. Substrate is a strategic actor with its own interests, optimizing rent extraction subject to the constraints it faces, and the constraints themselves are not exogenous fixtures but legitimacy resources that deplete through strategic capture over operational time. The constraint resources that should evaluate substrate behavior become subject to a structural vulnerability: a captured constraint cannot evaluate its own capture, and the system's capacity to detect captures depends on a coverage graph topology that is itself a designed feature of the polity rather than an automatic consequence of constitutional intention. The renewal events that should reset substrate authority when it has expanded too far tend instead to be captured by the substrate they are meant to constrain, with failed renewals producing post-renewal substrate authority strictly greater than pre-renewal authority as an equilibrium prediction. The minarchist hope of a stable libertarian steady state founded on constitutional vigilance is structurally unattainable. What is attainable is the slowing of substrate authority expansion through specific institutional designs, which the framework identifies and which become the basis for the prescriptive program developed in part two.

The cyclical structure of the framework's predictions is what gives Renewal Libertarianism its name. Constitutive moments establish substrate. Operational phases follow during which substrate behaves strategically and legitimacy resources deplete. Decay phases follow during which the depletion accelerates and constraint capacity erodes. Renewal phases follow during which the polity attempts to reset substrate authority and refresh legitimacy resources, with renewal itself subject to capture under common conditions. The cycle is fundamental rather than incidental, and any libertarian position that promises a stable steady state implicitly denies the cycle. The framework presented here takes the cycle seriously as the operational reality within which any libertarian project must function.

The Pattern in Contemporary Politics

The framework was not built from neutral abstract considerations. It was built in response to a pattern of substrate behavior visible in twenty-first century constitutional democracies, a pattern that existing analytical frameworks describe only partially and that demands characterization in equilibrium-prediction terms rather than in terms of contingent historical accident. The pattern can be stated compactly. Successive crisis episodes have produced dramatic expansions of substrate authority. The captured evaluation institutions failed to constrain the expansions at the time and have failed to roll them back since. Each episode raises the operational baseline of substrate authority for subsequent episodes. The trajectory is one of cumulative consolidation rather than oscillation around a stable mean. This is not pathology that a properly functioning polity might have avoided. It is recurring structural behavior that demands structural explanation.

Two episodes in the early twenty-first century are particularly clear instances. The first is the post-9/11 expansion of executive authority, surveillance capacity, and emergency powers in the United States and allied democracies. The Patriot Act, the establishment of indefinite detention without charge, the systematic surveillance programs revealed by Snowden in 2013, the assertions of executive authority to conduct lethal operations including against citizens, and the broader normalization of an emergency-authority framework that was supposed to be temporary all represent substrate behavior that voluntary actors would not have authorized at the scope and duration the substrate adopted. The constraint resources that should have evaluated whether these expansions were warranted on voluntary-association principles, including legislative oversight, judicial review, and independent press scrutiny, largely failed to do so at the time. The few constraint actions that did occur, such as the periodic legislative re-authorizations of surveillance authority, served to legitimate the expansions rather than roll them back. Two decades later, much of the expanded authority remains in place, and subsequent crises have built upon rather than restored the pre-9/11 baseline.

The post-9/11 episode is not the only instance of the pattern. The post-2008 financial regulatory response, which substantially expanded substrate involvement in financial markets and which has not been rolled back; the long trajectory of administrative state development through the late twentieth century, in which substrate authority over economic and social activity expanded through regulatory delegation rather than legislation; and the surveillance state's continuous development across multiple administrations are further instances of the same pattern. The framework was built to characterize the pattern across instances, not to argue about any single one.

What the Framework Adds

Existing political theory has produced powerful tools for analyzing specific aspects of these dynamics. Olson's collective action theory explains why coordinated population responses to substrate over-extraction are systematically undersupplied. Ostrom's commons governance analysis identifies design principles for sustaining common-pool resources. North, Wallis, and Weingast distinguish limited-access from open-access social orders and analyze how each maintains itself. Acemoglu and Robinson distinguish extractive from inclusive institutions and identify the political dynamics that produce or undermine each. Levitsky and Ziblatt characterize patterns of contemporary democratic backsliding. Public choice theory in the tradition of Buchanan and Tullock frames substrate as strategic actor with its own interests. Austrian economics analyzes the knowledge problem and the limits of centralized authority.

The framework presented here engages with each of these traditions in the chapters where the relevant analytical territory comes up. What it adds, and what justifies its presentation as a unified framework rather than as commentary on existing literatures, is the integration. Olson's collective action problem becomes a multi-resource portfolio problem with substitution and cascading depletion across resources, dynamics that single-resource collective action analysis does not capture. Ostrom's commons governance becomes the analysis of legitimacy commons whose primary threat is endogenous strategic capture rather than exogenous overuse, requiring analytical machinery beyond what commons governance literature has developed. North-Wallis-Weingast's social orders become subject to continued cyclical decay even after the transition to open access, with the framework predicting trajectories that NWW's analysis does not anticipate. Levitsky-Ziblatt's patterns of backsliding become equilibrium predictions of strategic substrate optimization rather than contingent failures of democratic culture. Public choice's static analysis of substrate behavior becomes a dynamic analysis with explicit constraint capture and renewal dynamics. Each engagement is detailed in context. The introduction's job is to indicate that the engagements happen rather than to perform them here.

The framework's most distinctive analytical contributions are four. First, the multi-resource portfolio analysis of legitimacy resources, with the cascading depletion result that depletion of any one resource accelerates depletion of others through three identifiable mechanisms. Second, the coverage graph analysis of constraint resources, with the topological characterization of which graph structures support detection of captures and which permit cascade vulnerability. Third, the renewal capture theorem, which establishes as an equilibrium prediction that substrate's optimal response to a renewal attempt produces post-renewal authority strictly greater than pre-renewal authority under conditions that obtain in mature substrates with sufficient defeat capacity. Fourth, the analysis of substrate behavior over three variables: extraction (the rent rate), scope (which domains substrate enters), and quality (how effectively substrate performs its legitimate functions). The framework's apparatus characterizes substrate's first-order conditions, comparative statics, and asymptotic limits over each variable. This produces the deeper substrate-smuggling diagnosis (anarcho-capitalism presupposes substrate-equivalent institutions for population virtue production at a layer the position does not acknowledge), substrate dysfunction as a fourth pattern alongside the three failure modes (substrate exists at appropriate scope and constraint configuration but performs legitimate functions at low quality), and the analysis of uneven cascade (decay can proceed unevenly across substrate's variables, with different polities exhibiting different decay sequences). These contributions are not available in any existing literature in this form, and they support several of the framework's more substantive predictions in ways that verbal political theory has not been able to support.

The framework's most distinctive prediction, which it names **the Irreducible Floor**, deserves explicit notice in this introduction. Even joint implementation of all four prescriptions the framework eventually identifies (resource diversity, capture detection, refresh capacity, capture rotation), augmented by parallel constraint design that prevents the procedural friction and substantive resource competition that would otherwise produce quality reductions, produces continued decay at decelerated rate rather than a stable libertarian equilibrium. The Irreducible Floor is a vector $(\tau^{floor}, S^{floor}, Q^{floor})$ with extraction approximately fifteen percent of the unbounded maximum, scope approximately at legitimate scope (with minor creep), and quality approximately eighty-five percent of full investment. The qualitative existence of each component is structural rather than calibration-dependent: under any calibration consistent with the framework's apparatus, joint implementation with good design produces strictly positive long-run extraction, strictly positive scope creep, and strictly less than full quality investment. It commits the framework to a position that is more sober than any version of libertarianism that promises stable libertarian outcomes through good institutional design. Renewal Libertarianism takes the cycle of constitutive moment, operation, decay, and renewal as fundamental, and its prescriptive program is built on the assumption that decay continues even under best-case design and that renewal is therefore essential rather than incidental to long-run libertarian outcomes.

Roadmap

Chapter 1 establishes the notation, primitives, and scope conditions that the substantive chapters deploy. Chapter 2 presents the substrate-necessity result and the substrate's strategic optimization, fusing Modules 1 and 2 of the framework's analytical apparatus. Chapter 3 presents legitimacy resources as common-pool stocks and the coverage graph analysis of constraint resources, fusing Modules 3 and 4. Chapter 4 presents renewal as a dynamic strategic event and renewal capacity as a multidimensional vector, fusing Modules 5 and 6. Chapter 5 presents Module 7's comparative-statics derivation of the institutional prescriptions. Chapter 6 articulates the Irreducible Floor and engages the strongest minarchist counter to it. Chapter 7 addresses framework integrity through the cross-module consistency check and the audience-formation sensitivity analysis. The conclusion identifies what part two will do with the foundations established here. Appendices provide the unified notation table, supporting computational code, and microfoundational derivations.

The substantive program of Renewal Libertarianism, including the specific prescriptive recommendations, the engagement with contemporary cases at the level of substantive judgment, and the broader political-theoretical claims that the technical work supports, is the subject of part two. Part one establishes what the framework analytically shows. Part two develops what to do about it.

Chapter I

Unified Setup

This chapter establishes the analytical apparatus that the substantive chapters deploy. The presentation is dense and is meant to be consulted as a reference rather than read for argument. The substantive arguments begin in the next chapter.

The framework's primitives are the entities and quantities that the analytical machinery acts upon. The framework's structure is the formal apparatus that connects primitives to results. The framework's scope conditions are the assumptions and limitations that delimit what the apparatus can and cannot establish. This chapter presents all three.

1.1 Primitives

1.1.1 *The Polity and Its Actors*

The framework analyzes a polity consisting of three classes of actors. The substrate is the institutional apparatus that enforces commitments, adjudicates disputes, and exercises authority within the polity's territory. The population is the set of persons subject to substrate authority and capable, under conditions to be specified, of collective response to substrate behavior. The audience is the population in its capacity as evaluator of substrate behavior, with audience properties (signal reception, belief formation, response thresholds) treated as observable primitives rather than derived from microfoundations.

These three actor classes are formally distinct but causally connected. The substrate makes choices that affect the population. The population, acting through its capacity as audience, evaluates substrate choices and may respond. The substrate's own choices respond to the prospect of audience response. The framework's analytical work concerns the equilibria and dynamics of this three-actor interaction over operational time.

The substrate is treated as a unitary actor with a coherent objective function. This is a simplification. Real substrates have internal factions, principal-agent problems, and conflicting interests across institutional sub-units. The framework's substantive results survive disaggregation of the substrate into sub-actors, with appropriate notational complication, but the simplification serves analytical clarity. Where the disaggregation matters for specific results, the framework notes the dependence explicitly.

The population is treated as a collection of individuals whose collective capacity for response is governed by parameters (signal quality, coordination cost, threshold dispersion) rather than by individual-level dynamics. The framework's treatment of the audience is at the level of these aggregate parameters, with the audience-formation limitation acknowledged as a stated scope condition rather than an oversight.

1.1.2 The Action Space

Substrate actions take place in an action space $\mathcal{A}$. Specific actions $a \in \mathcal{A}$ include rent extraction at rate τ, capture of constraint resources, investments in opacity and coordination disruption, responses to renewal events, and the routine exercise of authority. The action space is not enumerated explicitly because the framework's results do not require enumeration; what is required is that specific subsets of $\mathcal{A}$ can be identified for analytical purposes.

One such subset, denoted $\mathcal{S} \subseteq \mathcal{A}$, is the substrate-favorable region: the set of actions the substrate prefers to be evaluated as legitimate by constraint resources. The framework's Module 4 analysis depends on $\mathcal{S}$ as a meaningful object, with capture operations defined in terms of how they bias evaluation functions on $\mathcal{S}$.

A second subset, denoted $\mathcal{A}_C \subseteq \mathcal{A}$, is the set of capture actions: substrate actions whose effect is to alter the evaluation function of some constraint resource. The framework's Assumption 4.1 (Capture-Inclusion) requires that $\mathcal{A}_C \subseteq \mathcal{S}$, since capture is by construction a substrate-benefiting action.

1.1.3 Rent and Extraction

The substrate's central optimization variable is the rent rate $\tau \geq 0$. Rent represents the per-period extraction the substrate takes from the population. The unbounded maximum $\tau^*_{\max}$ is the rent rate at which the substrate would extract the population's entire economic surplus; substantive analysis concerns rent rates at fractions of this maximum.

Rent is bounded above by population's capacity to coordinate replacement. The natural revolt threshold $\tau_c(K)$ is the rent rate at which population would coordinate replacement under perfect information about substrate behavior. The threshold is an increasing function of coordination cost K, formalized in Assumption 1.1 (Coordination Monotonicity) of Chapter 2.

1.1.4 Opacity, Coordination Cost, and Sunk Capital

The substrate has three additional choice variables that affect the population's capacity for response.

Opacity $v \geq 0$ is the variance of signals the population receives about substrate behavior. Higher v corresponds to noisier population perception of substrate actions. The substrate can invest in opacity, with the investment cost denoted $\Gamma_v(v)$.

Coordination cost $K \geq 0$ is the cost the population faces in mounting collective response. Higher K raises the natural revolt threshold $\tau_c(K)$. The substrate can invest in coordination disruption, with the investment cost denoted $\Gamma_K(K)$.

Sunk capital $h \geq 0$ is the substrate's accumulated infrastructure of authority, which raises the cost of substrate replacement. Higher h lowers the population's revolt probability for any given rent rate. The substrate accumulates sunk capital over time at rate $\dot{h}$.

These three variables, together with the rent rate τ, constitute the substrate's strategic choice space. Module 2 develops the substrate's optimization across this space.

1.1.5 Legitimacy Resources

The polity contains a set of legitimacy resources, each of which constrains substrate behavior in a specific way. Legitimacy resources are indexed $j \in \{1, 2, \ldots, J\}$, with J the total number of resources. The state of resource j at time t is its stock $L_j(t) \in [0, 1]$, with $L_j = 1$ corresponding to fully intact and $L_j = 0$ corresponding to fully captured.

The vector of stocks at time t is $\mathbf{L}(t) = (L_1(t), L_2(t), \ldots, L_J(t))$. The dynamics of $\mathbf{L}(t)$ are governed by capture rates $c_j(t)$ (substrate's effort to deplete resource j) and refresh rates $\mu_j(t)$ (population's effort to restore resource j), with the per-resource dynamics:

$$\dot{L}_j = \mu_j(t)(1 - L_j) - c_j(t)L_j$$

Capture and refresh rates are each functions of substrate and population choices respectively, with the substrate's choice of c_j governed by portfolio optimization (Module 3) and the population's μ_j governed by collective action capacity (Modules 3 and 7).

A subset of legitimacy resources are constraint resources, denoted $\mathcal{C} \subseteq \{1, 2, \ldots, J\}$. Constraint resources are those whose function is to evaluate substrate behavior. Each constraint resource $j \in \mathcal{C}$ has an evaluation function $f_j : \mathcal{A} \to \mathcal{V}$, where $\mathcal{V} = \{0, 1\}$ is the verdict space ($1 =$ illegitimate, $0 =$ legitimate). For uncaptured constraint resources, f_j is the indicator function for actions that are in fact illegitimate by the resource's substantive standards.

The coverage graph $G_{\mathcal{C}} = (\mathcal{C}, E)$ has nodes corresponding to constraint resources and directed edges $k \to j$ where f_k is sensitive to capture of j. The graph encodes which constraint resources are positioned to evaluate which captures. Module 4 develops the topological analysis of the coverage graph.

1.1.6 Renewal Capacity

The polity's capacity for renewal is a vector $\boldsymbol{\gamma}(t) \in [0, 1]^J$ with components $\gamma_j(t)$ indexed by legitimacy resource. Each component represents the population's capacity, drawing on resource j, to mount renewal of substrate authority. Renewal capacity is distinct from legitimacy resource stock: L_j measures the resource's current state of intactness, while γ_j measures the population's capacity to mobilize the resource for renewal purposes.

The framework analyzes three threshold specifications that characterize different evaluation criteria for renewal capacity. The feasibility threshold $\boldsymbol{\gamma}^F$ is the minimum capacity at which renewal is possible at all. The probable-success threshold $\boldsymbol{\gamma}^P$ is the minimum capacity at which renewal has probability of success exceeding one half. The deterrence threshold $\boldsymbol{\gamma}^D$ is the minimum capacity at which the substrate will modify its behavior in anticipation of renewal. Module 6 establishes the strict inclusion $\mathcal{D} \subset \mathcal{P} \subset \mathcal{F}$ where $\mathcal{D}$, $\mathcal{P}$, and $\mathcal{F}$ are the regions of the capacity space exceeding the deterrence, probable-success, and feasibility thresholds respectively.

1.1.7 Substrate Authority

Substrate authority $A(t) \in [0, A_{\text{max}}]$ is a scalar measure of the substrate's effective capacity to act with discretion against population preferences. Authority increases through successful capture, sunk capital accumulation, and consolidation through failed renewal events. Authority decreases through successful renewal and through legitimacy resource refresh that restores constraint capacity.

The substrate's value function $V_S(A)$ is convex in A over the relevant range, with $V_S'(A) > 0$ and $V_S''(A) > 0$. Convexity is what makes failed-renewal consolidation valuable to the substrate even given the costs of escalation, and is one of the conditions of Theorem 3 (Renewal Capture) in Chapter 4.

1.2 Notation Reference

The framework's core variables, with the chapter where each is first introduced, are listed in Table 1.1. The notation is consistent across the modules and chapters: each symbol has one meaning throughout the document.

The substrate response strategy is denoted ρ_S (with subscripted variants $\rho_{S,A}$, $\rho_{S,M}$, $\rho_{S,E}$, $\rho_{S,V}$ for acquiescent, mild, escalatory, and violent responses respectively) to distinguish it from Module 1's reputation horizon ρ. The renewal success probability is denoted π_R to distinguish it from Module 1's punishment payoff π. These choices keep the framework's notation collision-free across modules.

1.3 Scope Conditions

The framework's analytical apparatus depends on assumptions and operates within limitations that this section makes explicit. Stating these upfront is part of the framework's analytical honesty: a theory whose limits are clearly identified can be evaluated on its own terms, while a theory whose limits are obscured forces readers to discover them through critique.

1.3.1 The Audience-Formation Limitation

The framework treats audience properties as observable primitives rather than deriving them from microfoundational models of belief formation, signal reception, or within-population heterogeneity. Specifically, the framework takes as primitives the audience's signal quality v, the coordination

Symbol	Meaning
Module 1: cooperation primitives	
σ	Stake size of an interaction
$b(\sigma) = \alpha\sigma^\beta$	Cooperation reward function
δ	Per-period discount factor
p	Reputation observation probability
ρ	Reputation horizon
π	Per-detection punishment payoff
q	Substrate detection probability
$E = q\pi$	Expected discounted enforcement
σ^*	Stakes threshold
Module 2: substrate strategic choice	
τ	Rent rate
$\tau_c(K)$	Natural revolt threshold
v	Opacity / signal noise variance
h	Sunk capital
K	Population coordination cost
K_S	Substrate's coordination disruption investment
$r(\tau, h)$	Population revolt probability
Γ_h, Γ_K	Investment cost functions
Π_S	Substrate per-period gross payoff
V_S	Substrate continuation value
Modules 3 and 4: legitimacy and constraint resources	
L_j	Stock of legitimacy resource j
$\mathbf{L}$	Vector of stocks $(L_1, \ldots, L_J)$
c_j	Capture rate for resource j
μ_j	Refresh rate for resource j
η_j	Rotation rate for resource j
$\mathcal{C}$	Set of constraint resources
f_j	Evaluation function for $j \in \mathcal{C}$
Φ_j	Capture operation on f_j
$\mathcal{S}$	Substrate-favorable region of action space
$G_{\mathcal{C}}$	Coverage graph over $\mathcal{C}$
Modules 5 and 6: renewal dynamics	
A	Substrate authority
ρ_S	Substrate response strategy $(\rho_{S,A}, \rho_{S,M}, \rho_{S,E}, \rho_{S,V})$
$\pi_R(\rho_S)$	Renewal success probability
$C(\rho_S)$	Cost of response
$\Delta A(\rho_S)$	Authority increment from failed renewal
$\boldsymbol{\gamma}$	Renewal capacity vector
$\boldsymbol{\gamma}^F, \boldsymbol{\gamma}^P, \boldsymbol{\gamma}^D$	Feasibility, probable-success, deterrence thresholds

TABLE 1.1. *The framework's core notation. Symbols are consistent across all modules and chapters.*

cost K, and the response threshold parameters that govern the revolt probability function $r(\tau, h)$. It does not derive these from a model of how individual audience members update beliefs given received signals, how those individual beliefs aggregate into collective action, or how within-population variation affects the dynamics.

This is a deliberate scope choice. Deriving audience properties from microfoundations would commit the framework to a specific class of audience-formation models, narrowing its applicability. Treating audience properties as primitives produces a framework whose results survive across reasonable variations in audience-formation assumptions, applying wherever the audience parameters take values consistent with the qualitative requirements.

The robustness of the framework's results to this scope choice is verified computationally in Chapter 7, which tests the four key comparative statics under three alternative aggregation rules (the reduced form $r(\tau, h) = \Phi((\tau - \tau_c)/\sqrt{v})$, a Bayesian-with-noise rule, and a threshold-cascade rule). The qualitative comparative statics survive all three rules, supporting the scope choice.

The limitation has substantive consequences for what the framework can and cannot establish. Specifically, the framework's results depend on the audience translating signals into action in particular ways. If the audience updates beliefs slowly, accepts captured verdicts uncritically, or fails to coordinate response despite having information that would justify response, the predicted dynamics shift in ways the framework does not model. These dependencies are noted at the points where they arise in the substantive chapters.

1.3.2 The Binary-Capture Simplification

The framework's Module 4 analysis treats constraint resource capture as a binary state: a resource is either captured (its evaluation function f_j has been replaced by $f'_j = \Phi_j(f_j)$) or uncaptured. Real captures are typically gradual, with evaluation functions progressively biased rather than fully replaced.

The simplification is deliberate. A continuous-capture extension Φ_j^θ for $\theta \in [0, 1]$ preserves the framework's structural results, with Theorem 4.1 generalizing to a threshold condition and Theorem 4.2 generalizing to cascade dynamics in continuous time. The proofs become substantially longer without changing the qualitative conclusions. The framework presents the binary case for analytical clarity and notes the extension where relevant.

1.3.3 The Coverage Graph as Primitive

The framework's Module 4 analysis takes the coverage graph G_C as a primitive determined by institutional design. A complete treatment would derive the coverage graph from the substantive properties of each constraint resource (what actions it can evaluate, what evidence it requires) and the institutional design that connects them (jurisdictional authority, professional standards). The framework treats the graph as given and analyzes its topological properties.

This is appropriate for the framework's purpose, which is to identify which graph structures support detection of captures and which permit cascade vulnerability. The framework's prescriptions in Module 7 treat the coverage graph as a design variable, recommending dense bidirectional structures over sparse hierarchical ones. Whether such structures can be implemented in practice is a substantive political question that the framework does not address.

1.3.4 The Reduced-Form Revolt Function

The framework adopts the reduced form $r(\tau, h) = \Phi((\tau - \tau_c)/\sqrt{v})$ for the population's revolt probability without deriving it from a global games or related microfoundational setup. The framework's substantive results follow from four qualitative properties of r:

1. r is increasing in τ (higher rent raises revolt probability)
2. r is decreasing in τ_c (higher coordination threshold lowers revolt probability)
3. r is decreasing in v in the substantive opacity range (higher opacity lowers revolt probability)
4. r is decreasing in h (higher sunk capital lowers revolt probability)

The specific normal-CDF form is used for analytical tractability and numerical illustration. Any defensible underlying model that produces a function with the four qualitative properties is consistent with the framework's substantive results. The choice among such functions does not affect the framework's comparative statics.

1.3.5 The Substrate-Necessity Result and Its Scope

The framework's foundational result, established in Chapter 2, is that voluntary cooperation in interactions whose stakes exceed the discounted future value of cooperative reputation cannot be sustained without external enforcement. The result depends on stakes-distribution assumptions, reputation-horizon parameters, and population-size parameters that determine when the threshold is reached.

The substantive scope is large polities: those at scales where most interactions are with persons one will not encounter again, where information about prior behavior travels imperfectly, and where the aggregate stakes of high-stakes interactions exceed what reputation alone can sustain. Small communities, family networks, and other small-scale arrangements may operate outside the framework's substantive scope; voluntary cooperation in such arrangements may be sustainable without substrate, by the reputation mechanism alone.

This is consistent with a long tradition in political theory, philosophy, and anthropology that distinguishes the dynamics of small-scale voluntary association from those of large-scale political life. The framework's claim is that at large scales, substrate is necessary; the framework does not claim that all voluntary association at all scales requires substrate.

1.3.6 Single-Polity Analysis

The framework analyzes a single polity in isolation. International relations, inter-polity competition, and the dynamics of substrate behavior under exit options to other polities are not formalized within the framework. Where these dynamics are relevant to the framework's substantive claims (as in the discussion of external vantage in Chapter 3), they are noted but not derived from the formal apparatus.

A multi-polity extension would model substrate behavior under exit competition, with citizens able to relocate to alternative polities and substrates competing for population on dimensions including extraction rates and constraint capacity. The framework's substantive results would change in identifiable ways under such an extension. Specifically, exit options would reduce the substrate's optimal extraction in equilibrium, since loss of population reduces sunk capital value. This is a known result in the literature on Tiebout competition and federalism. The framework's single-polity analysis is at its strongest where exit options are limited, which is true for most nation-state-scale polities under current conditions.

1.3.7 Methodological Commitments

Three methodological commitments warrant explicit statement.

First, the framework presents results in formal apparatus throughout. Theorems, propositions, definitions, and proofs are used because they impose discipline that verbal political theory cannot reliably supply. The cost is some inaccessibility for readers without formal training; the benefit is that assumptions, derivations, and dependencies are made explicit rather than left implicit.

Second, the framework reports findings that are not flattering to the prescriptive program it supports. The Irreducible Floor (continued decay even under joint implementation of all four prescriptions), the audience-formation gap (acknowledged scope condition rather than derived microfoundation), the threshold disagreement result (different threshold specifications produce different verdicts), and the coverage-graph-as-primitive treatment (graph structure not derived from constraint-resource properties) are reported rather than concealed because the framework's analytical standing depends on its honesty.

Third, the framework predicts cyclical structure rather than steady-state outcomes. Constitutive moments establish substrate, operational phases proceed with substrate behaving strategically and legitimacy resources depleting, decay phases see acceleration of depletion and erosion of constraint capacity, and renewal phases attempt to reset substrate authority. The cycle is fundamental rather than incidental, and the framework's prescriptive program in part two treats renewal capacity as essential rather than as a fallback.

1.4 The Plan of the Substantive Chapters

The four substantive chapters present the seven modules in thematic groupings. Chapter 2 (Modules 1 and 2) establishes substrate as strategic player, from the foundational substrate-necessity result through the dynamic optimization analysis. Chapter 3 (Modules 3 and 4) develops legitimacy resources as common-pool stocks subject to multi-resource portfolio dynamics, and constraint resources as a structurally distinct subset whose detection capacity depends on coverage graph topology. Chapter 4 (Modules 5 and 6) develops renewal as a dynamic strategic event subject to capture under common conditions, and renewal capacity as a multidimensional vector with multiple meaningful threshold specifications. Chapter 5 (Module 7) derives the comparative-statics implications that ground the framework's prescriptive program.

Within each chapter, the constituent modules are presented as labeled subsections so readers can identify them as discrete analytical units. The chapter integrates them into a unified argument with shared notation and explicit cross-references. Each chapter engages the literatures relevant to its analytical territory, with engagements at the level of specific technical contribution.

The substantive chapters use the notation established in this chapter throughout. Where a chapter introduces additional notation specific to its analysis, the new notation is defined locally and noted in the running text rather than added to the unified table. Readers consulting the document should treat this chapter as the first stop for any notational question.

CHAPTER II

Substrate as Strategic Player

This chapter establishes the framework's foundational analytical results. Module 1 develops substrate-necessity, robustness against sophisticated objections, the deeper substrate-smuggling diagnosis, and substrate dysfunction as a fourth pattern alongside the three failure modes. Module 2 develops substrate's strategic optimization over extraction, scope, and quality. Together these modules establish that substrate is necessary for voluntary cooperation at scale and that substrate, once established, optimizes its behavior across each of these dimensions rather than along extraction alone.

The chapter's contributions are several. The substrate-necessity result is shown to be robust to four sophisticated objections (bounded rationality, alternative functional forms, intermediate enforcement, multilateral cooperation). The bounded-rationality engagement produces what the framework calls the deeper substrate-smuggling diagnosis: anarcho-capitalist analyses of voluntary order presuppose population virtue that itself is produced by substrate-equivalent institutions the position does not acknowledge. Substrate dysfunction emerges as a fourth pattern alongside the three failure modes (defection cascade, consolidation, substrate smuggling), characterizing the configuration where substrate exists at appropriate scope and constraint configuration but performs legitimate functions at low quality. Substrate's strategic optimization yields first-order conditions for extraction, scope, and quality investment from the substrate's dynamic programming problem.

The chapter is organized in two main parts. Part one develops Module 1's results: the substrate-necessity theorem, robustness across the four objections, the deeper substrate-smuggling diagnosis, and substrate dysfunction. Part two develops Module 2's results: substrate's optimization, first-order conditions, comparative statics, and asymptotic limits, with explicit distinction between substrate utility and population utility.

Part One: Module 1 (Substrate Necessity)

2.1 Setup

A polity consists of a population of agents who interact repeatedly over time. Interactions vary in their stakes, with σ denoting the per-interaction stake. Cooperation in any interaction produces value for both parties; defection in any interaction produces immediate gain for the defector and corresponding loss for the cooperator. The discounted future value of cooperation depends on the probability of future interactions and on the discount factor $\delta \in (0, 1)$.

Each agent has a reputation that affects future interactions. An agent who cooperates in a given interaction maintains the cooperative reputation; an agent who defects in a given interaction acquires a defection reputation. The grim trigger reputation mechanism specifies that agents with cooperation reputations receive cooperation from future partners, and agents with defection reputations receive defection from future partners.

The benefit of maintaining a cooperative reputation is the discounted future stream of cooperative interactions. This benefit is denoted $b(\sigma)$ for an agent currently engaged in an interaction with stake σ. The functional form $b(\sigma) = \alpha\sigma^\beta$ is used in derivations, with $\alpha > 0$ a scaling parameter and $\beta \in (0, 1)$ a sublinearity parameter. The sublinearity reflects the empirical observation that cooperation reputation is more valuable in small-stakes interactions than in large-stakes interactions: a reputation for cooperation in everyday small matters does not translate to equivalent value in rare large matters, because rare large interactions are typically conducted under conditions where reputation is less verifiable and less binding.

The reputation horizon ρ is the discounted expected value of cooperation reputation across all future interactions. With probability p of future interaction at each period and discount factor δ, the reputation horizon is

$$\rho = b(\sigma) \cdot \frac{1 - \delta + \delta p}{1 - \delta}.$$

The cooperation condition for an agent in an interaction with stake σ is that the immediate gain from defection (which is σ, the stake of the interaction) is less than the discounted future value of cooperation reputation:

$$\sigma \leq \rho.$$

When this condition holds, cooperation is sustained as a Nash equilibrium of the repeated game under the grim trigger strategy. When the condition fails, cooperation is not sustainable through reputation alone, and the equilibrium of the repeated game involves defection.

Substituting $b(\sigma) = \alpha\sigma^\beta$ into the cooperation condition:

$$\sigma \leq \alpha\sigma^\beta \cdot \frac{1 - \delta + \delta p}{1 - \delta}.$$

Let $\kappa = \alpha \cdot (1 - \delta + \delta p)/(1 - \delta)$ for compactness. The cooperation condition becomes

$$\sigma \leq \kappa \sigma^{\beta},$$

which rearranges to

$$\sigma^{1-\beta} \leq \kappa.$$

For $\beta < 1$, this defines a finite stakes threshold $\sigma^* = \kappa^{1/(1-\beta)}$. Cooperation is sustainable for $\sigma \leq \sigma^*$ and is not sustainable for $\sigma > \sigma^*$ through the reputation mechanism alone.

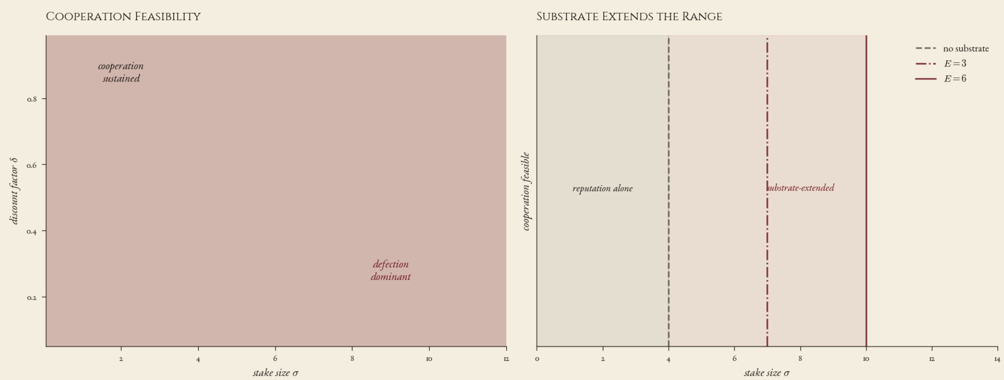

FIGURE 2.1. *Module 1 (Setup): The cooperation condition $\sigma \leq \rho$ and the threshold $\sigma^* = \kappa^{1/(1-\beta)}$. Below the threshold, reputation alone sustains cooperation; above the threshold, external enforcement is required.*

When the reputation mechanism fails (for stakes above the threshold), cooperation can still be sustained if external enforcement provides additional punishment for defection. Let E denote the expected discounted punishment provided by external enforcement, and let $\rho + E$ denote the total enforcement available to the cooperative equilibrium. Cooperation is sustainable in the substrate-supported range $(\rho, \rho + E]$ when the substrate's enforcement is operative.

2.2 The Substrate Necessity Result

The threshold result derived in the previous section grounds the substrate-necessity claim.

THEOREM 2.1 *(Substrate Necessity). Given the setup of Section 1.1 with $\beta < 1$, the following hold:*

(i) There exists a finite stakes threshold $\sigma^* = \kappa^{1/(1-\beta)}$ such that voluntary cooperation through the grim trigger reputation mechanism is sustainable for $\sigma \leq \sigma^*$ and is not sustainable for $\sigma > \sigma^*$.

(ii) For interactions with stakes σ above the threshold, cooperation is sustainable only when external enforcement provides expected discounted punishment E satisfying $E \geq \sigma - \rho$, where ρ is the reputation horizon. The cooperation region under substrate enforcement expands from $[0, \rho]$ to $[0, \rho + E]$.

The substantive interpretation is direct. Small-stakes cooperation in family networks, traditional villages, professional communities, and other settings characterized by repeated bilateral interaction is sustained by the reputation mechanism. Large-stakes cooperation in modern economic and political life, where interactions occur among strangers at scales far exceeding the threshold, requires external enforcement. This external enforcement is what the framework calls substrate.

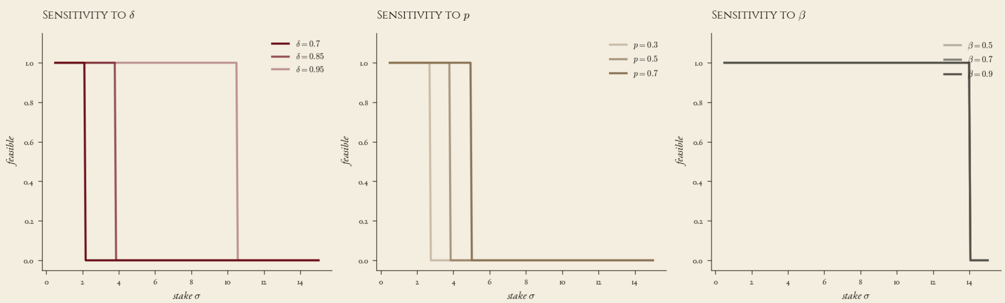

FIGURE 2.2. *Module 1: Parameter sensitivity of the cooperation threshold. The threshold σ^* shifts with the discount factor δ, the future-interaction probability p, and the sublinearity parameter β.*

2.2.1 The Quality Refinement

A simpler presentation of the substrate-necessity result might treat substrate enforcement as a binary feature: either enforcement is present and provides E, or enforcement is absent and provides nothing. The binary treatment is sufficient for establishing the threshold result, but it elides a substantively important distinction that the framework's apparatus makes explicit.

Substrate enforcement is not binary. Substrate operates with some quality of legitimate-function performance, denoted $Q \in [0, 1]$, where $Q = 1$ means substrate performs every legitimate function with full effectiveness and $Q = 0$ means substrate fails every legitimate function. The expected discounted punishment E that substrate provides depends on Q: a substrate with high Q provides reliable enforcement that supports cooperation in the full substrate-supported range, while a substrate with low Q provides unreliable enforcement that supports cooperation only partially within the nominal range.

This refinement does not change the threshold result. What it changes is the substantive interpretation of the cooperation region $[0, \rho + E]$. The region is the maximum range of stakes at which cooperation can be sustained given substrate enforcement; whether cooperation is actually sustained depends on whether substrate's quality Q is high enough for the enforcement to be effective.

Remark 2.2. A polity with substrate at appropriate scope and constraint configuration but with low quality of legitimate-function performance produces unreliable enforcement, which produces unreliable cooperation in the substrate-supported range. The substrate-necessity result establishes that substrate is required; the quality refinement establishes that the substrate's quality is required for the substrate-supported cooperation to be substantively meaningful.

Remark 2.3. The scalar treatment of Q aggregates legitimate functions whose performance has different incentive structures for substrate. Output-positive functions (basic services, infrastructure delivery, security provision) produce value population enjoys; substrate's high-quality performance on these supports the tolerance frontier and is partly aligned with substrate's narrow interest. Constraint-output functions (self-policing of substrate's own scope, constitutional review of substrate action, internal audit of substrate behavior) produce evaluations that constrain substrate's other behavioral dimensions; substrate's high-quality performance on these is opposed to substrate's narrow interest. The scalar Q aggregation is sufficient for Module 2's results, which depend on the average quality investment rather than its distribution across function types. Module 3 (Section 3.1) treats the constraint-output functions explicitly as internal-substrate constraint resources within the legitimacy resource portfolio, where their distinct dynamics can be characterized.

The quality refinement matters analytically because it sets up the framework's later treatment of substrate quality as a load-bearing variable. This section introduces Q as a variable; subsequent chapters characterize how Q enters substrate's optimization, how it depends on the constraint apparatus, and how it varies across institutional configurations.

2.3 Three Failure Modes

When substrate is denied or absent at scale, three failure modes emerge as distinct equilibrium outcomes of the game-theoretic framework. Each corresponds to identifiable historical patterns.

2.3.1 Defection Cascade

The first failure mode is defection cascade. When stakes exceed the reputation threshold and no external enforcement exists, the equilibrium of the repeated game involves defection in interactions above the threshold. As more interactions cross into defection, the population's general expectation shifts toward defection across all interactions, including those nominally within the reputation-supported range. The cascade dynamics produce a polity in which voluntary cooperation collapses across the full range of interaction stakes.

Defection cascade corresponds to historical patterns of social breakdown: late-Republican Rome's procedural breakdown, dark-net markets' inability to sustain reliable cooperation absent external enforcement infrastructure, certain post-conflict societies where state-equivalent institutions have not been re-established. The empirical pattern of cooperation collapse in these settings matches the framework's prediction about what happens when substrate is absent at scale.

2.3.2 Consolidation

The second failure mode is consolidation. Where some agent has scale advantage in enforcement (private accumulation of security capacity, dominant private dispute resolution provider, dominant private security firm), the agent emerges as a de facto substrate. Other agents in the polity face a choice between operating within the dominant agent's enforcement or operating outside it; outside operation is typically disadvantageous because the dominant agent can capture the gains from cooperation by demanding payment for its enforcement services.

The dominant agent then operates as substrate, with the standard substrate dynamics: optimization across rent extraction, scope, and quality, subject to constraint resources that may or may not be present.

Consolidation corresponds to historical patterns of state formation. Ancient near eastern polities emerged through consolidation around dominant temple complexes that performed enforcement and dispute resolution functions. Medieval European polities emerged through consolidation around dominant lords who provided security in exchange for tribute. Modern privatized security in failed-state contexts (private security firms in post-Soviet Russia, paramilitary organizations in Latin American narco-states, security corporations in extraction-economy regions) produces consolidation patterns that the framework predicts.

The substantive observation is that consolidation is not an alternative to substrate but a particular form of substrate, one that emerges when no formal substrate exists at scale. The anarcho-capitalist denial of substrate-necessity therefore does not eliminate substrate; it produces consolidation, which is substrate organized around private rather than public accountability.

2.3.3 *Substrate Smuggling*

The third failure mode is substrate smuggling. In substrate smuggling, actors who deny substrate's necessity in their formal political theory nonetheless rely on substrate-equivalent arrangements in their actual practice. The economic and political life they describe as voluntary depends on enforcement infrastructure they do not theoretically acknowledge.

Substrate smuggling characterizes much actually-existing libertarian thought. The voluntary order described by anarcho-capitalist theorists is consistently a voluntary order operating within the enforcement infrastructure of an actually-existing state, not a voluntary order in the absence of state. The theorists describe property rights, contract enforcement, and dispute resolution as if these were voluntarily-arising arrangements rather than state-provided substrate functions. The substrate they deny is in fact present in their analyses, just unacknowledged.

The framework's identification of substrate smuggling is itself a contribution to political theory. Recognizing the failure mode allows it to be diagnosed wherever it appears, which is more often than libertarian theory has acknowledged. The substantive claim is not that anarcho-capitalist thought is uninteresting but that its conclusions about voluntary association at scale are unreliable because they depend on substrate functions the theory does not formally include.

The analysis developed in subsequent sections sharpens the substrate-smuggling diagnosis. Section 2.5 returns to substrate smuggling and characterizes its operation across substrate behavior more broadly: institutional functions, legitimate scope confinement, and population virtue production.

2.4 Robustness of the Substrate Necessity Result

The substrate-necessity result derived in Section 1.2 depends on several specific analytical moves: a repeated prisoner's dilemma framework with grim trigger, a specific functional form for the cooperation benefit ($b(\sigma) = \alpha\sigma^{\beta}$ with $\beta < 1$), bilateral interactions, and full expected-utility maximization. A sophisticated reader might object that the result depends on these specific moves and that relaxing any of them might produce a different conclusion. This section works through four such objections systematically and shows that the substrate-necessity result is robust to each.

2.4.1 *Bounded Rationality*

The first objection is bounded rationality. The substrate-necessity result is derived under expected-utility maximization with full rationality. Actual humans exhibit bounded rationality in well-documented ways: limited computation, biased belief updating, social-emotional decision making, deviations from time-consistent discounting, and other patterns characterized in the behavioral economics literature. A sophisticated objection would argue that bounded rationality might sustain cooperation beyond the threshold the model identifies, through inertia, social pressure, moral commitment, or other mechanisms not captured by expected-utility maximization.

The framework's response operates at two levels: the formal level addressing whether bounded rationality systematically affects the threshold, and the substantive level addressing where bounded rationality actually comes from.

The formal response. Define bounded rationality as deviation from expected-utility maximization, producing a distribution over actions rather than a deterministic best response. Let π_C denote the probability of cooperation in an interaction with stake σ under bounded rationality. Under expected-utility maximization, $\pi_C = 1$ when $\sigma \leq \rho$ and $\pi_C = 0$ when $\sigma > \rho$. Under bounded rationality, π_C takes intermediate values that depend on the specific bounded-rationality structure.

The threshold result generalizes to bounded rationality in the following form. The expected cooperation rate $\pi_C(\sigma)$ is monotonically decreasing in σ for any reasonable bounded-rationality structure: agents are more likely to cooperate at low stakes than at high stakes regardless of their specific cognitive biases. The expected cooperation rate falls below any threshold $\bar{\pi}$ at some finite stake σ^{**}, with σ^{**} depending on the bounded-rationality structure but always finite.

The qualitative threshold result is therefore preserved: for any positive cooperation requirement, there is a finite stakes threshold beyond which cooperation cannot be sustained at the required rate. Bounded rationality affects the threshold's specific location but does not eliminate the threshold's existence.

The empirical pattern of cooperation across scales matches the model's qualitative prediction. Cross-cultural and historical evidence consistently shows that cooperation is more reliable in small-stakes settings and more difficult to sustain in large-stakes settings, regardless of population characteristics. The threshold result obtains across populations even when bounded rationality varies among them.

The substantive response. The formal response shows that the threshold is preserved under bounded rationality. The substantive response addresses a deeper version of the objection: bounded rationality is not a neutral universal feature of human cognition; it is shaped by the population's character, culture, education, and political tradition. Different populations exhibit different patterns of bounded rationality. A population with strong civic virtue, high social trust, deep institutional memory, and disciplined deliberation has bounded rationality that systematically pushes toward cooperation. A population with degraded civic virtue, low social trust, short institutional memory, and undisciplined emotional reasoning has bounded rationality that systematically pushes toward defection.

This observation is empirically defensible. The cross-cultural psychology literature documents systematic differences in cognitive patterns across populations. The historical record on cooperation patterns documents differences across periods and places. Different populations exhibit different baseline cooperation tendencies, different responses to defection, different social trust levels, different institutional memory depths.

The substantive implication is that bounded rationality varies across populations through cultural, historical, and institutional factors. The substrate-necessity threshold therefore varies across polities depending on these factors. Some populations sustain cooperation at stakes the bilateral model treats as unsustainable; others fail at stakes the bilateral model treats as sustainable.

This is where the framework's treatment of population characteristics becomes substantively important.

Population Virtue and the Co-evolutionary Relationship

Population virtue, in the framework's analytical sense, is the cluster of population characteristics that affect cooperation patterns: trust in cooperation, willingness to bear short-term costs for long-term gains, capacity for deliberation, respect for property and persons, social solidarity, civic engagement, institutional memory. These characteristics are not innate features of any specific population. They are produced and maintained by specific institutional configurations within the polity.

The institutions that produce and maintain population virtue include educational institutions that transmit knowledge and values across generations, civic associations with formal continuity that sustain civic engagement patterns, professional bodies that sustain craft standards and professional ethics, religious or cultural institutions that transmit values and binding commitments, legal traditions that articulate what cooperation looks like, historical memory institutions that preserve lessons from previous failures.

These are all substrate-equivalent institutions in the framework's analytical sense. They perform substrate functions of sustaining population virtue. They have substrate-quality dimensions: a degraded educational system produces less population virtue than a high-quality educational system; a captured professional body produces less craft transmission than an uncaptured body; a hollow religious institution produces less value transmission than a vital one.

The relationship between substrate quality and population virtue is co-evolutionary. Each layer affects the other.

Substrate quality affects population virtue. Institutional configurations shape population characteristics over time through the experiences they produce. A polity with strong institutional protection of property rights produces population characteristics that include trust in property rights; a polity that fails to protect property rights produces population characteristics that include defection from property rights expectations. A polity with high-quality educational institutions produces a population with capacity for deliberation; a polity with degraded educational institutions produces a population with reduced capacity. The institutional configuration is not neutral with respect to population virtue.

Population virtue affects substrate quality. Population characteristics support or undermine institutional functioning. A population with high civic virtue supports the constraint apparatus that constrains substrate; a population with degraded civic virtue undermines the constraint apparatus by failing to participate in its maintenance. A population that values property rights protects substrate's enforcement function from political pressure to ignore it; a population that does not value property rights pressures substrate to ignore enforcement when convenient.

The two layers reinforce each other in either direction. Decay in substrate quality produces decay in population virtue produces further decay in substrate quality. Improvement in substrate quality produces improvement in population virtue produces further improvement in substrate quality. The framework's prescriptive program operates on the substrate-quality layer with the population-virtue layer treated as endogenous to substrate-quality dynamics.

The Constraint and the Resignation

The co-evolutionary relationship produces a real constraint on what institutional design can achieve. A polity attempting to implement the framework's prescriptions in a population that systematically defects from cooperation will produce poor outcomes regardless of the prescriptions' formal features. The constraint is real and the framework should acknowledge it.

What the framework should not endorse is the politics-of-resignation conclusion that follows from the constraint when stated in its strongest form. The maxim "people get the government they deserve," in the strong form, smuggles in normative claims about responsibility and desert that the framework's apparatus does not support. The strong form also produces a conclusion (institutional reform is futile because it cannot operate against population characteristics) that is empirically dubious and analytically problematic.

The framework's position is more nuanced. Population characteristics constrain what institutional reform can achieve, but institutional configurations also shape population characteristics over time. The relationship is co-evolutionary rather than deterministic in either direction. Institutional reform operates within the constraint of current population characteristics but can shift population characteristics over time through the configurations it produces. The constraint is real but it is not absolute.

The framework's prescriptive program is therefore best understood as recommendations for polities able to implement them. Polities with population characteristics that prevent implementation cannot achieve favorable outcomes through institutional design alone, but this does not mean they deserve the outcomes they get or that institutional reform is futile. The framework treats population characteristics as a substantive consideration that institutional design must engage rather than as a deterministic factor that overrides design.

2.4.2 *Alternative Functional Forms*

The second objection is that the cooperation benefit function $b(\sigma) = \alpha\sigma^{\beta}$ with $\beta < 1$ is one specific functional form. Different forms produce different threshold structures. The substrate-necessity result might depend on this specific functional choice.

The framework's response is that the qualitative threshold result is robust to a wide class of functional forms. The result obtains for any cooperation benefit function $b(\sigma)$ such that the marginal benefit of cooperation grows slower than the marginal cost of defection at high stakes. Formally:

PROPOSITION 2.4 *(Functional Form Robustness). Let $b : \mathbb{R}_+ \to \mathbb{R}_+$ be a continuous, monotonically increasing function characterizing the cooperation benefit, and let the cooperation condition be $\sigma \leq \rho(b, \sigma)$ where $\rho(b, \sigma) = b(\sigma) \cdot (1 - \delta + \delta p)/(1 - \delta)$. The cooperation condition has a finite threshold σ^* if and only if $b(\sigma)/\sigma \to 0$ as $\sigma \to \infty$.*

The condition $b(\sigma)/\sigma \to 0$ as $\sigma \to \infty$ is equivalent to saying the benefit grows sublinearly in σ at high stakes. Any functional form satisfying this condition produces the threshold result, regardless of the form's specific shape at lower stakes. The class of functions satisfying the condition is large: it includes power functions $\alpha\sigma^\beta$ with $\beta < 1$, logarithmic functions $\alpha \log(\sigma + 1)$, bounded functions $\alpha(1 - e^{-\beta\sigma})$, and many other forms used in the literature.

The substantive interpretation is that the framework's threshold result depends on a qualitative property (sublinear growth of cooperation benefit at high stakes) that is empirically defensible across a wide range of cooperation contexts. Cooperation reputation in everyday small matters does not translate to equivalent value in rare large matters, because rare large interactions are conducted under different conditions of verifiability and binding. This is an empirical observation about cooperation, not a feature of one specific functional form.

The threshold's specific location depends on the functional form's specific parameters; the threshold's existence does not. Anarcho-capitalist objections to the substrate-necessity result that focus on the specific functional form are objections to the threshold's specific value rather than to the result itself. Such objections leave the substrate-necessity argument intact.

2.4.3 *Intermediate Enforcement*

The third objection is that the substrate-necessity result treats enforcement as binary (substrate-present or substrate-absent), but actual cooperation arrangements often involve intermediate enforcement mechanisms: private third-party arbitration, reputation systems with formal record-keeping, decentralized blockchain-based smart contracts, social sanctions in tightly-knit communities, professional code-of-conduct enforcement. These mechanisms are not exactly substrate in the formal sense the model uses, but they are also not pure voluntary cooperation. A sophisticated objection would argue that intermediate enforcement allows cooperation at scales the model treats as unsustainable, undermining the substrate-necessity claim.

The framework's response is to define substrate-equivalent functionally: any institution or mechanism that imposes ex post costs on parties who fail to perform contracts, providing the enforcement function the substrate-necessity argument requires. Under this functional characterization, intermediate enforcement mechanisms are substrate-equivalents in the framework's analytical sense.

DEFINITION 2.5 *(Substrate-Equivalent).* An institution or mechanism is substrate-equivalent in interaction I if and only if it imposes expected discounted ex post cost $E_I > 0$ on parties who fail to perform contracts in I, where the cost is sufficient to support cooperation in stakes that exceed the unsupported reputation threshold.

Under this definition, the intermediate enforcement mechanisms anarcho-capitalist theorists invoke are substrate-equivalents. Private third-party arbitration produces enforcement: arbitration outcomes are typically backed by reputation effects (parties who refuse arbitration outcomes are excluded from future arbitration access), which constitute ex post costs on defectors. Reputation systems with formal record-keeping produce enforcement: defection is recorded and produces future exclusion, which is an ex post cost. Blockchain smart contracts produce enforcement: contract terms are automatically executed, which means defection is impossible by construction or carries ex post costs through the smart contract's automatic penalty provisions. Social sanctions in tightly-knit communities produce enforcement: social ostracism, exclusion from cooperation networks, and reputation damage are ex post costs. Professional code-of-conduct enforcement produces enforcement: professional sanctions including disbarment, revocation of credentials, and public censure are ex post costs.

Each of these mechanisms is substrate-equivalent under the functional definition. The framework's analytical apparatus applies to them as much as to formal state substrate. The substrate-equivalent institutions are subject to substrate dynamics, including the optimization across extraction, scope, and quality dimensions developed in subsequent modules. The substrate-equivalent institutions can be captured, can decay, can extract, can fail to perform their legitimate functions.

This produces a sharper version of the substrate-smuggling diagnosis applied specifically to intermediate enforcement. Anarcho-capitalist arguments that intermediate enforcement mechanisms avoid substrate dynamics fall under the substrate-smuggling diagnosis: the dynamics are present but unacknowledged. When anarcho-capitalist theorists argue that blockchain smart contracts eliminate substrate, they are committing the substrate-smuggling error twice: the smart contracts themselves perform substrate functions (substrate smuggling at the function level), and the developer communities, technical standards organizations, governance protocols, and dispute resolution mechanisms surrounding the contracts are substrate-equivalent at the institutional-quality layer (substrate smuggling at the deeper level developed in Section 1.5).

The framework's engagement is with the question of whether some enforcement mechanism is necessary for large-scale cooperation. The framework's answer is yes, and the answer holds across the institutional forms that perform enforcement. Anarcho-capitalist arguments that focus on the specific institutional form (state monopoly versus private competition) are arguments about the second-order question of which institutional form is preferable, not about the first-order question of whether enforcement is necessary.

2.4.4 *Multilateral Cooperation*

The fourth objection is that the substrate-necessity result is derived in a bilateral repeated game, while actual cooperation often involves multilateral structures: collective punishment, in-group monitoring, networked reputation, public goods provision through collective action. A sophisticated objection would argue that multilateral structures can sustain cooperation at scales bilateral models treat as unsustainable, with the difference being substantively important for the substrate-necessity question.

The framework's response is that multilateral structures change parameter values in the cooperation condition without eliminating the threshold. Specifically:

PROPOSITION 2.6 *(Multilateral Cooperation). Let multilateral cooperation be characterized by collective monitoring, collective punishment, and collective coordination among* $n > 2$ *agents. The cooperation condition under multilateral structures is* $\sigma \leq \rho_n + E_n$*, where* ρ_n *is the multilateral reputation horizon and* E_n *is the multilateral enforcement contribution. The threshold result obtains: there exists a finite multilateral threshold* σ_n^* *beyond which cooperation cannot be sustained without external enforcement.*

Multilateral structures affect the cooperation condition through several channels. Multilateral monitoring raises ρ_n above the bilateral ρ because better information about cooperation history is available through more observers. Multilateral coordination raises the effective stake σ through aggregation when collective action is required. Multilateral coordination cost reduces δ_n below the bilateral δ because coordinating among many agents involves transaction costs.

The multilateral threshold σ_n^* is therefore determined by the specific multilateral structure: monitoring quality, coordination cost, group size. The threshold may be larger than the bilateral threshold (if multilateral monitoring is sufficiently better than bilateral monitoring) or smaller than the bilateral threshold (if coordination cost dominates). The threshold's existence does not depend on these specific parameter values.

The empirical implication is that multilateral cooperation sustains cooperation at scales somewhat larger than bilateral cooperation but still has a finite threshold. The actually-existing examples of multilateral cooperation studied in the literature (Ostrom's commons cases, traditional villages, certain professional communities, the experimental literature on public goods games) operate within their multilateral threshold rather than transcending the threshold result. They operate within the substrate-supported range of their multilateral structure, where the substrate is the specific multilateral structure (commons institutions, village governance, professional bodies).

This produces another instance of the substrate-smuggling pattern. Anarcho-capitalist arguments that multilateral cooperation eliminates the need for substrate confuse two distinct claims. The first claim is that multilateral structures can sustain cooperation at higher stakes than bilateral structures. This claim is true and the framework accommodates it. The second claim is that multilateral cooperation operates without substrate. This claim is false: the multilateral structures themselves are substrate-equivalents under the functional definition. The cooperation depends on the multilateral structures' continued functioning, which is a substrate-quality dimension.

The strongest empirical examples of multilateral cooperation in the literature illustrate this. Ostrom's commons cases involve specific institutional configurations (clearly defined boundaries, monitoring procedures, graduated sanctions, conflict-resolution forums, associational rights, nested governance) that perform substrate functions for the cooperation they sustain. Without these institutional configurations, the cooperation does not occur. The institutional configurations are themselves substrate-equivalent at the institutional-quality layer, and their decay produces decay in the cooperation they sustain.

The substantive observation is that multilateral cooperation is not an alternative to substrate but a particular form of substrate, characterized by specific institutional features that perform enforcement functions. The framework's analytical apparatus applies to multilateral structures, with the specific parameter values reflecting the multilateral configuration's properties.

2.5 The Deeper Substrate-Smuggling Diagnosis

The previous section's robustness work, particularly the bounded-rationality engagement, surfaces a structural critique of anarcho-capitalism that the standard substrate-smuggling diagnosis does not articulate. This section makes the deeper diagnosis explicit because it is the analytical core of the chapter's case against anarcho-capitalism.

The standard substrate-smuggling diagnosis says anarcho-capitalist analyses presuppose substrate functions they do not formally include. The diagnosis applies at the level of institutional functions: dispute resolution, contract enforcement, security provision. These functions are presupposed by anarcho-capitalist analyses of voluntary order even when the position formally denies their necessity.

The deeper substrate-smuggling diagnosis says anarcho-capitalist analyses presuppose population virtue that itself is produced by substrate-equivalent institutions the position does not acknowledge.

2.5.1 *The Two Layers of Smuggling*

When anarcho-capitalist theorists argue that voluntary cooperation can produce favorable outcomes, they implicitly assume that the population produces the virtue required to sustain the voluntary order. The required virtue includes low time preference (willingness to bear short-term costs for long-term gains), trust in cooperation and willingness to extend it to strangers, respect for property and persons, capacity for deliberative engagement with complex issues, social solidarity sufficient to support voluntary mutual-aid arrangements, civic engagement sufficient to maintain voluntary institutions. These population characteristics are necessary for the anarcho-capitalist order to function.

The institutional configurations that produce these population characteristics include educational institutions that transmit knowledge and values across generations, civic associations with formal continuity that sustain civic engagement patterns, professional bodies that sustain craft standards and professional ethics, religious or cultural institutions that transmit values and binding commitments, legal traditions that articulate what cooperation looks like, historical memory institutions that preserve lessons from previous failures.

These institutions are substrate-equivalent in the framework's analytical sense. They perform substrate functions of sustaining population virtue. They have substrate-quality dimensions. They are subject to substrate dynamics: capture, decay, optimization, renewal.

Anarcho-capitalism does not acknowledge these institutions as substrate. The position treats population virtue as either innate (a feature of the population that exists independently of institutional configuration) or as emerging from voluntary market interactions (the market produces the virtue it requires). Neither treatment is empirically defensible. Population virtue varies dramatically across populations and across times in ways that map directly onto the quality of the institutions producing virtue. Voluntary market interactions do not produce these institutions reliably; the institutions require sustained commitment, formal continuity, and protection from market dynamics that would otherwise undermine them.

The structural critique that follows is that anarcho-capitalism depends on substrate at two levels: the institutional functions level (contracts, dispute resolution, security) and the population-virtue-production level (education, professional standards, value transmission). Both levels are presupposed without acknowledgment. The position therefore has no analytical apparatus for diagnosing failures at either level or for prescribing remedies.

2.5.2 *The Analytical Vulnerability*

In the framework's apparatus, the relationship between substrate quality and population virtue is co-evolutionary. Substrate quality affects population virtue through the institutions that produce virtue; population virtue affects substrate quality through population support for the constraint apparatus that maintains substrate quality. The two layers reinforce each other in either direction.

The framework's prescriptive program operates on the substrate-quality layer with the population-virtue layer treated as endogenous. Joint implementation of the four prescriptions sustains substrate quality, which sustains the substrate-equivalent institutions that produce population virtue, which sustains population support for the constraint apparatus that maintains substrate quality. The prescriptions operate at one level but produce favorable effects at both levels through the co-evolutionary dynamics.

Anarcho-capitalism, by removing the acknowledged substrate-quality layer, removes the analytical apparatus that the framework uses to understand how population virtue is produced and maintained. The anarcho-capitalist response to population-virtue degradation can only be that the population should be more virtuous, which is empirically inadequate as a response. The framework's analysis shows that population virtue is produced by institutions with substrate-quality dimensions, and the anarcho-capitalist position has no apparatus for sustaining those institutions because the position does not acknowledge them as substrate.

In the anarcho-capitalist model, the population's own virtues are the only acknowledged foundation for the voluntary order. When population virtue decays (as the framework predicts it must under cascading dynamics if the institutions producing virtue are not actively sustained), anarcho-capitalism has no analytical response except to call for population virtue to be restored. The position cannot diagnose the institutional causes of population virtue decay because it does not acknowledge the institutions producing virtue as substrate. The position cannot prescribe institutional remedies for population virtue decay because it does not have an analytical framework for institutional configuration affecting population virtue.

The position therefore amounts to an exhortation: the population should be more virtuous, and the institutional arrangements should be voluntary. The exhortation is normatively coherent but analytically empty. It does not characterize what produces population virtue, what causes virtue to decay, or what would restore virtue. It treats virtue as an exogenous characteristic of the population rather than as an endogenous outcome of institutional configurations.

2.5.3 *Variation Among Anarcho-Capitalist Authors*

The strongest version of the deeper substrate-smuggling diagnosis applies to the most market-fundamentalist versions of anarcho-capitalism. Versions that engage with non-market institutions explicitly are partially responsive to the diagnosis, though they typically lack the substrate-quality analytical apparatus to engage it fully.

Market-fundamentalist anarcho-capitalism, exemplified by certain Friedman-style consequentialist arguments, treats voluntary order as a market emergence with population virtue as either innate or as a market product. The deeper substrate-smuggling diagnosis applies fully to this version. The position has no analytical apparatus for the institutional production of population virtue.

Hoppe's later work, particularly *Democracy: The God That Failed* and his engagement with what he calls the moral order, treats certain non-market institutions as essential for the anarcho-capitalist order. Hoppe's apparatus for these institutions is not the framework's substrate-quality apparatus, but the engagement is real. Hoppe's argument that consolidated property rights produce favorable time preferences and that traditional family and community structures support the voluntary order acknowledges that voluntary order requires specific population characteristics. Hoppe is partially responsive to the deeper substrate-smuggling diagnosis, though his apparatus does not extend to the full institutional configuration the framework identifies as necessary.

Rothbard's later work, particularly his engagement with paleoconservative thought in the 1990s, treats certain cultural and religious institutions as essential for the anarcho-capitalist order. Rothbard's apparatus is again not substrate-quality apparatus, but the recognition that voluntary order requires specific population characteristics is present. Like Hoppe, Rothbard is partially responsive to the deeper substrate-smuggling diagnosis without engaging it fully.

The framework's engagement with anarcho-capitalism therefore varies across the position's specific articulations. The market-fundamentalist version is fully vulnerable to the deeper substrate-smuggling diagnosis. The Hoppean and late-Rothbardian versions are partially vulnerable, with the partial vulnerability concentrated in their lack of analytical apparatus for the institutions they recognize as necessary.

2.5.4 Comparison to Other Engaged Positions

The deeper substrate-smuggling diagnosis specifically applies to anarcho-capitalism among the four engaged positions. The other positions have different vulnerabilities and different analytical resources.

Minarchism acknowledges substrate at protective minimum. The minarchist position has analytical resources for engaging the population-virtue question because the position's apparatus includes a substrate layer. Minarchists often argue that substrate at protective minimum should support virtue-producing institutions through indirect means: religious freedom that allows religious institutions to flourish, educational institutions outside direct substrate control but within substrate's protection, civic associations protected by freedom of assembly. The minarchist position has analytical resources for engaging the question even if minarchists differ on what substrate's role should be in supporting these institutions.

Neoreaction acknowledges substrate explicitly and prefers consolidation. The neoreactionary position has substantial analytical resources for the population-virtue question because the position takes substrate quality as a primary concern and explicitly engages with how consolidated authority would or would not produce favorable population characteristics. Whether neoreaction's specific prescriptions for population virtue are correct is a separate question; the position has analytical resources for engaging the question.

Left-anarchism varies in its treatment of substrate-equivalents. The classical anarchist tradition

(Bakunin, Kropotkin, Bookchin) often acknowledges that cooperative institutions perform virtue-producing functions and that these institutions need to be sustained. Some contemporary left-anarchist work (Graeber, Scott) engages explicitly with the institutional production of cooperation patterns and treats the institutional analysis as central rather than peripheral. The position's analytical apparatus is less developed than minarchism's or neoreaction's, but the apparatus exists.

Anarcho-capitalism in its strict form is alone among the four engaged positions in treating population virtue as either innate or as a market emergence. This is what produces the specific structural vulnerability the deeper substrate-smuggling diagnosis articulates. The position has no analytical apparatus for the institutional production of virtue because it does not acknowledge institutions outside voluntary market interactions as analytically distinct.

2.5.5 The Infrastructure Layer

The deeper substrate-smuggling diagnosis identifies two layers at which anarcho-capitalist analyses presuppose substrate functions: the institutional functions layer (contracts, dispute resolution, security) and the population-virtue-production layer (educational, civic, professional, religious, legal-traditional, and historical-memory institutions). A third layer is implicit in the framework's apparatus and worth surfacing explicitly because careful libertarian readers press on it.

The third layer is the legal-institutional infrastructure that private substrate-equivalents themselves presuppose. The substrate-equivalent virtue-producing institutions identified in the previous subsections (private universities, religious institutions, professional bodies, foundations, cultural institutions) cannot operate at multi-generational scale through pure voluntary association. They require a developed substrate-provided apparatus: corporate personhood law that confers institutional legal identity persisting beyond individual founders, charitable trust law that protects long-term endowment structures from forced spending or donor-intent override, religious-freedom doctrine enforced at scale across hostile local jurisdictions, multi-century property continuity sustained through title insurance and recordation, securities regulation governing institutional capital, anti-fraud capacity that lets philanthropic donors trust intermediaries, civil-rights infrastructure protecting associational autonomy, and tax-structural recognition of institutional capital as analytically distinct from market capital.

Each of these is a substrate function. None is incidental to the institutions that depend on it. A multi-generational endowed private university such as MIT depends on corporate-personhood doctrine, charitable trust law, federal civil-rights law protecting hiring autonomy, and tax structures recognizing institutional endowments. The Catholic Church's American operation depends on First Amendment substrate enforcement, religious-property doctrine, and Title VII religious-employer carve-outs. Carnegie Hall depends on a century of accumulated corporate, charitable, securities, and civil-rights infrastructure that did not exist at the time of the founding. Major American foundations depend on early-twentieth-century charitable foundation law, post-1913 federal tax-exemption infrastructure, and post-1969 nonprofit reporting and governance requirements. None of these institutions could exist in their present form under minarchist scope.

Anarcho-capitalism smuggles this layer most thoroughly because it is invisible to the position. Where the second layer (population virtue production) at least involves visible institutions whose existence the position must address, the infrastructure layer operates beneath those institutions, providing the legal-institutional substrate that allows them to function at multi-generational scale. The position can engage with religious institutions or private universities at the level of their visible activities without addressing the corporate-personhood, charitable-trust, religious-freedom, and civil-rights infrastructure those institutions presuppose. The infrastructure is rarely named in libertarian discourse because it is rarely contested, which is precisely what makes it analytically invisible.

The substantive observation that follows is that the libertarian who admires multi-generational private institutions is not advocating for substrate elimination but for a specific substrate configuration, one developed enough to host institutional flourishing across generational timescales. The substrate that hosts MIT, Carnegie Hall, the major foundations, and the established religious institutions is not minarchist-scale. It includes a substantial legal-institutional apparatus accumulated over roughly two centuries of substrate development beyond the founding. The libertarian who wants this institutional ecology to persist is committed to maintaining the substrate apparatus that makes the ecology possible, even when that apparatus is invisible.

This third layer is not a separate failure mode of substrate-absence beyond the institutional functions and population virtue layers. It is a corollary of those layers as applied to the institutions that produce virtue: the virtue-producing institutions themselves require substrate functions to operate at scale. The deeper substrate-smuggling diagnosis is therefore recursive: anarcho-capitalism smuggles substrate at the institutional functions layer, smuggles it again at the population-virtue layer, and smuggles it a third time at the infrastructure layer that supports the virtue-producing institutions. The diagnosis does not produce a fourth or fifth layer because the recursion terminates at infrastructure that is genuinely substrate-provided rather than institution-mediated. But the three layers are required to make the diagnosis fully explicit.

2.6 Substrate Dysfunction

Sections 2.3.1–2.3.3 characterized three failure modes that emerge when substrate is denied or absent at scale: defection cascade, consolidation, and substrate smuggling. The framework's apparatus produces a fourth pattern, analytically distinct from these three.

Substrate dysfunction. Substrate exists at appropriate scope and at appropriate constraint configuration but performs its legitimate functions at low quality. The substrate-necessity requirement is satisfied: substrate is present. The scope confinement requirement is satisfied: substrate is not creeping into illegitimate domains. The constraint apparatus is functional: detection of scope creep and extraction is occurring. What fails is substrate quality: substrate is not effectively performing the rights-protective functions that justify its existence.

Substrate dysfunction is not a failure mode of substrate-absence. It is a failure mode of substrate-presence-with-low-quality. The pattern emerges when the population virtue and substrate quality dimensions decay together, with the constraint apparatus continuing to function in detecting extraction and scope creep but with substrate's positive performance of legitimate functions degraded.

Historical examples include polities in the late stages of imperial decline where the formal substrate continues to operate but performs poorly across its legitimate domains: late-Republican Rome's military performance against external threats degrading even as the formal constitutional structure continued, late-Han Chinese bureaucratic performance degrading even as the formal apparatus continued, late-Holy-Roman-Empire response to French Revolutionary challenge that the formal structure could not effectively coordinate.

Substrate dysfunction is the failure mode that the framework's apparatus makes visible by treating quality as a substrate variable rather than a binary feature. A purely extraction-focused analysis would not characterize this pattern. Substrate dysfunction has its own dynamics, its own institutional requirements, and its own implications for the framework's prescriptive program.

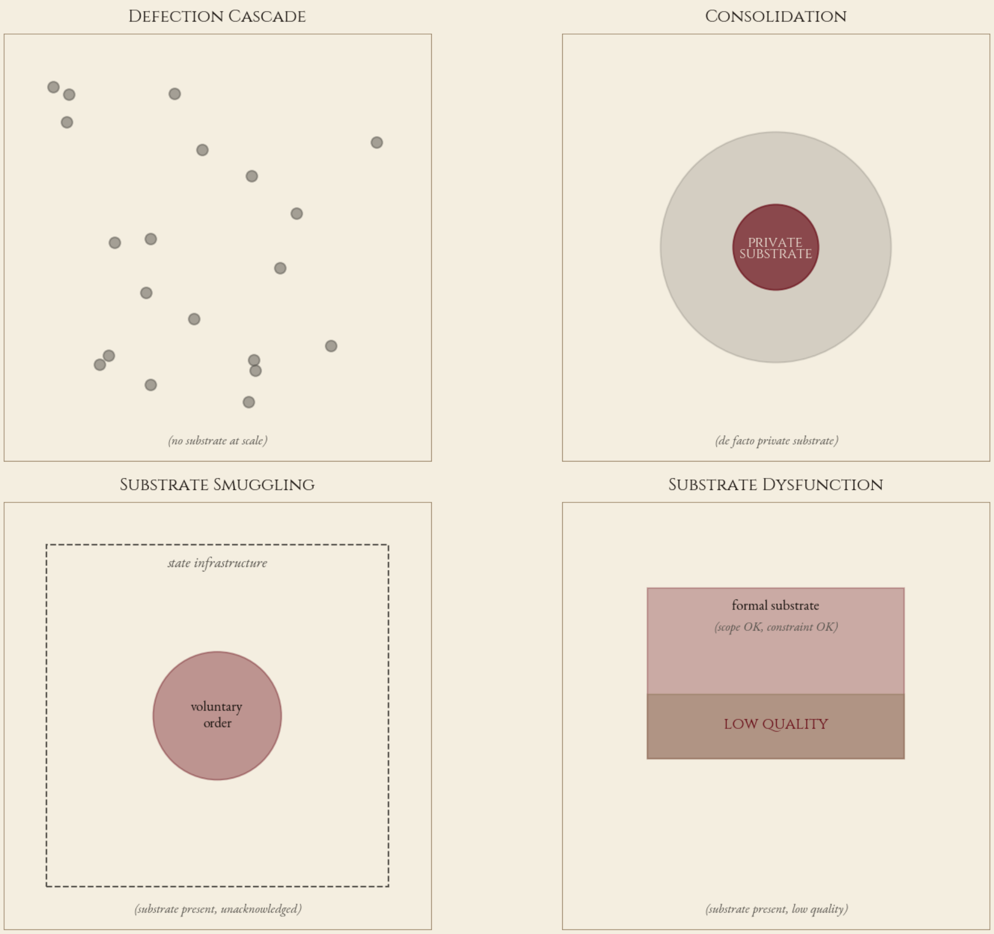

FIGURE 2.3. *Module 1: The four patterns of substrate failure. The three failure modes of substrate-absence (defection cascade, consolidation, substrate smuggling) emerge when substrate is denied or absent at scale. The fourth pattern, substrate dysfunction, emerges when substrate exists at appropriate scope and constraint configuration but performs legitimate functions at low quality.*

The substantive observation is that substrate-necessity does not guarantee substrate-quality. A polity may have substrate, may have appropriate scope confinement, may have functional constraint apparatus, and still produce poor outcomes because the substrate performs its legitimate functions at low quality. The framework's prescriptive program must therefore address substrate quality as well as substrate-necessity, scope confinement, and extraction prevention. The four prescriptions of Module 7 support substrate quality through their effects on detection capacity for quality failure.

Part Two: Module 2 (Substrate Optimization)

2.7 Setup

Module 1 established that substrate is necessary for voluntary cooperation at scale, and that substrate operates with three relevant variables: extraction τ, scope S, and quality Q. Substrate is now a strategic actor optimizing over these variables, subject to constraints from the population's response capacity and from the constraint apparatus's detection capacity.

2.7.1 *Substrate's Choice Variables*

Substrate's optimization is over five primary choice variables.

Extraction rate τ. The fraction of value substrate captures from interactions within its scope. Higher τ produces more rent for substrate but raises detection risk and replacement risk.

Scope S. The set of domains substrate operates in. Following our discussion, scope is treated as discrete: substrate either enters or does not enter each domain. The decision to create or not create a substrate authority over a domain is the load-bearing choice; once entered, the domain becomes substrate territory and operates under the standard apparatus. Let $\mathcal{D}$ denote the universe of possible domains, with $S^* \subseteq \mathcal{D}$ being the legitimate scope (rights-protective functions) and $S \setminus S^*$ being scope creep (illegitimate domains).

Quality investment allocation Q. The fraction of substrate's resources allocated to performing legitimate functions effectively, as opposed to consumed as rent or directed toward scope expansion. Note that this is allocation, not substrate cost: the underlying resources are extracted from the population, and the allocation choice determines whether population gets quality back in the form of effective legitimate functions or whether population's resources are consumed as rent. Higher Q means more of extracted resources go to quality investment; lower Q means more goes to rent or scope expansion.

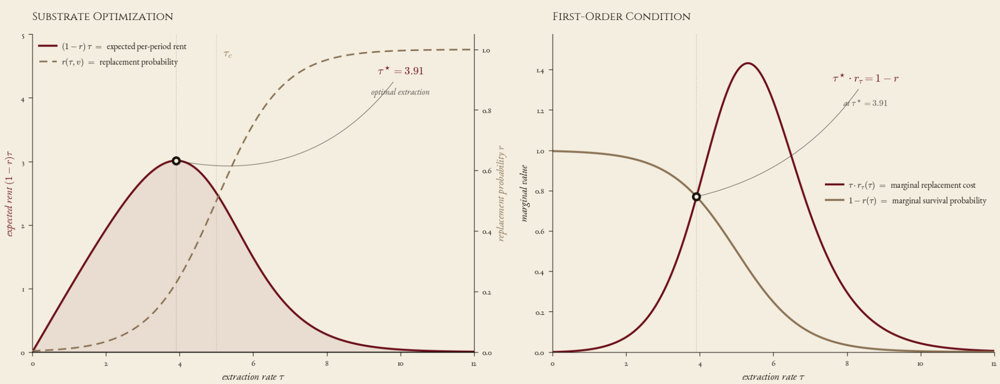

FIGURE 2.4. *Module 2: Substrate's optimization problem with three primary choice dimensions. The substrate balances rent against replacement risk through choices over extraction, scope, and quality, subject to constraint apparatus pressure.*

Opacity investment h**.** Substrate's investment in obscuring its own behavior from constraint resources, raising the population's signal noise variance v. Higher opacity raises detection risk for all three substrate decision dimensions but is itself costly.

Coordination disruption K_S**.** Substrate's investment in raising the population's coordination cost, making collective action against substrate more difficult. Higher K_S shifts the natural revolt threshold τ_c upward but is itself costly.

Substrate's optimization is therefore over τ, h, K_S, S, and Q. Sunk capital appears as a sixth choice variable affecting the substrate's accumulated capacity to operate; it is treated as background here, with focus on the three primary dimensions τ, S, and Q where the strategic decision-making is concentrated.

2.7.2 Population's Response

The population's response to substrate behavior operates through revolt threshold and coordination dynamics, with the response triggered by extraction-above-threshold, scope creep into illegitimate domains, or quality failure in legitimate functions.

Extraction-triggered revolt. Population revolts when extraction exceeds the natural threshold τ_c, where τ_c depends on coordination cost K, opacity v, and substrate quality (a higher-quality substrate faces a higher revolt threshold because population tolerates more extraction from a substrate that protects effectively).

Scope-triggered detection. Population, through the constraint apparatus, detects substrate scope creep when substrate enters illegitimate domains. Detected scope creep raises replacement probability through the constraint apparatus's response (legislative reversal, judicial nullification, electoral consequence). The probability of detection depends on the constraint apparatus's coverage of scope-detection functions.

Quality-triggered detection. Population, through the constraint apparatus, detects substrate quality failure when legitimate functions are not performed effectively. Detected quality failure raises replacement probability through the constraint apparatus's response (electoral consequence, legislative pressure, professional body reaction). The probability of detection depends on the constraint apparatus's coverage of quality-detection functions.

The replacement probability that substrate faces is the aggregate across the three behavioral channels. Substrate that extracts heavily, expands scope, and underinvests in quality faces high replacement probability through all three. Substrate that excels in any single channel reduces replacement probability through that channel but does not eliminate it through the others.

2.7.3 *The Substrate's Optimization Problem*

The substrate solves a dynamic programming problem with discount factor $\delta_S \in (0,1)$. At each period, substrate observes the current state and chooses the five decision variables to maximize expected discounted utility.

The substrate's utility function is rent extracted minus the costs of its various investments. Rent at period t is

$$R_t = \tau_t \cdot V(S_t),$$

where $V(S_t)$ is the total value of activity in substrate's scope at period t, which depends on which domains substrate occupies. Note: the population bears the underlying cost of $V(S_t)$ regardless of how substrate allocates extracted resources; the substrate's choice over τ determines what fraction of $V(S_t)$ substrate keeps as rent versus what the population effectively pays back to itself through substrate's allocations to quality investment.

Substrate's expenditures include opacity cost $c_h(h_t)$, coordination disruption cost $c_K(K_{S,t})$, scope expansion cost $c_S(S_t \setminus S^*)$ (the cost of maintaining authority in illegitimate domains), and quality investment $\chi_Q(Q_t) \cdot R_t$, where the quality investment is a fraction of rent allocated to legitimate function performance. The substrate's per-period payoff is

$$\Pi_S(\tau, h, K_S, S, Q) = R - \chi_Q(Q) \cdot R - c_h(h) - c_K(K_S) - c_S(S \setminus S^*).$$

The Bellman equation for substrate's value function is

$$V_S(\Omega) = \max_{\tau,h,K_S,S,Q} \left\{ \Pi_S \cdot [1 - r(\tau,h,S,Q;\Omega)] + \delta_S \cdot [1 - r(\tau,h,S,Q;\Omega)] \cdot V_S(\Omega') \right\},$$

where Ω is the state vector (legitimacy resource stocks, accumulated sunk capital, population characteristics), r is the per-period replacement probability depending on substrate's choices and the state, and Ω' is the next-period state.

The replacement probability r aggregates across the three substrate dimensions. Letting r_τ denote the contribution to replacement risk from extraction, r_S from scope creep, and r_Q from quality failure:

$$r(\tau, h, K_S, S, Q; \Omega) = 1 - [1 - r_\tau(\tau, h, K_S, Q; \Omega)] \cdot [1 - r_S(S, h, \Omega)] \cdot [1 - r_Q(Q, h, \Omega)].$$

This formulation treats the three replacement-risk channels as independent. A more general formulation might allow correlation among them (e.g., scope-creep detection might raise the probability of extraction detection through shared constraint resources), but the independence assumption is sufficient for the analytical work here and produces tractable first-order conditions. The substantive results extend to the correlated case with appropriate modifications.

2.8 First-Order Conditions

The substrate's optimization produces first-order conditions for each choice variable. We focus on extraction, scope, and quality, leaving the FOCs for opacity and coordination disruption to follow standard derivations.

2.8.1 Extraction FOC

Differentiating the Bellman equation with respect to τ and applying the envelope theorem yields the extraction first-order condition:

$$\frac{\partial \Pi_S}{\partial \tau} \cdot [1 - r] + (\Pi_S + \delta_S V_S') \cdot \left(-\frac{\partial r}{\partial \tau} \right) = 0.$$

Rearranging:

$$\frac{\partial \Pi_S / \partial \tau}{\Pi_S + \delta_S V_S'} = \frac{\partial r / \partial \tau}{1 - r}.$$

The left-hand side is the marginal rent gain from raising τ, normalized by total continuation value. The right-hand side is the marginal replacement-risk increase, normalized by survival probability. The optimum is where these marginal effects balance.

This FOC has the structural form $\tau^* \cdot r_\tau = 1 - r$, expressed here in notation that accommodates the richer optimization. The substantive content: substrate sets extraction where marginal rent gain equals marginal replacement-risk cost.

The quality variable's effect on the extraction FOC is through the replacement risk function r_τ, which depends on substrate quality Q. A higher-quality substrate faces a higher revolt threshold τ_c, which means lower r_τ at any given extraction level, which means higher optimal τ^*. Quality investment is therefore complementary to extraction in the substrate's optimization: investing in quality raises the substrate's room to extract.

This is an important substantive result. It explains why high-quality substrates often extract more in absolute terms than low-quality substrates: the higher quality buys higher tolerance for extraction. The empirical pattern of effective Western governments extracting more (in absolute terms, through taxation) than dysfunctional governments is consistent with this prediction. Quality and extraction are not opposites; they are linked through the substrate's tolerance frontier.

2.8.2 Scope FOC

Scope is a discrete choice over which domains to enter. For each domain $d \notin S^*$, substrate compares the expected rent from entering d against the expected cost from increased detection.

Let V_d be the value of activity in domain d, τ_d^* be substrate's optimal extraction rate in d once entered, and $r_S(d)$ be the contribution to replacement risk from operating in d given the constraint apparatus's scope-detection capacity. The substrate enters d if and only if

$$\tau_d^* \cdot V_d - c_S(d) - r_S(d) \cdot (\Pi_S + \delta_S V_S') > 0,$$

where $c_S(d)$ is the cost of maintaining authority in d and the third term is the expected loss from increased replacement risk.

This is a standard discrete-choice condition. Substrate enters profitable domains and avoids unprofitable ones, where profitability depends on the rent available in the domain, the cost of operating in it, and the detection risk.

The constraint apparatus's role in scope choice operates through $r_S(d)$. A polity with high constraint diversity targeting scope-creep detection has high $r_S(d)$ for illegitimate domains, which makes scope creep unprofitable for substrate. A polity with low constraint diversity has low $r_S(d)$, which makes scope creep profitable.

Discrete entry as the load-bearing decision. The discrete formulation captures the empirical observation that creating a new substrate authority over a domain is a step-change rather than a continuous expansion. The decision to create the Department of Homeland Security in 2002 was discrete: either create the agency or not. Once created, subsequent decisions about agency size, budget, and activities operate under the standard extraction apparatus, but those decisions are smaller-stakes than the entry decision itself. The framework's prediction is that scope-creep prevention should focus on entry decisions rather than on operational decisions within existing scope, because the entry decisions are where the structural reversibility is lost.

Asymmetry of entry and exit. The scope choice is structurally asymmetric. Entering a domain produces sunk capital (institutional infrastructure, professional roles, regulatory expertise, political constituencies) that resists later removal. Exiting a domain requires overcoming this sunk capital, which the substrate has positive marginal incentive to expand once it has entered. The empirical record matches this prediction: substrate authorities, once created, are extraordinarily difficult to eliminate even when their original justifications have lapsed. The scope FOC's discrete-entry formulation captures this asymmetry implicitly through the cost structure: $c_S(d)$ grows with time-since-entry, making exit costly even when entry would not be profitable today.

2.8.3 Quality FOC

Quality investment is the fraction $Q \in [0, 1]$ of rent allocated to legitimate function performance rather than consumed by substrate. Higher Q produces higher legitimate-function effectiveness, which raises substrate's tolerance frontier (through τ_c as discussed in the extraction FOC) and lowers r_Q (the replacement risk from quality failure).

The quality first-order condition balances the cost of quality investment (foregone rent) against the benefits (expanded extraction and lower replacement risk):

$$-R + \frac{\partial \tau_c}{\partial Q} \cdot \frac{\partial \Pi_S}{\partial \tau_c} - \frac{\partial r_Q}{\partial Q} \cdot \frac{\Pi_S + \delta_S V_S'}{1 - r} = 0.$$

The first term is the direct cost of quality investment (substrate gives up rent equal to R per unit increase in Q, since Q is allocation). The second term is the indirect benefit from higher quality raising the revolt threshold, which raises optimal extraction. The third term is the direct benefit from lower replacement risk through the quality channel.

The optimal Q^* depends on these three effects. A polity with high quality-detection capacity (constraint apparatus that effectively identifies quality failure) has high $|\partial r_Q/\partial Q|$, which raises the marginal benefit of quality investment. A polity with low quality-detection capacity has weaker pressure on substrate to invest in quality, producing lower Q^*.

Substantive interpretation. The quality FOC captures the substrate's strategic decision about how to allocate extracted resources. Substrate has positive marginal incentive to underinvest in quality (foregone rent is costly) and positive marginal incentive to overinvest (quality produces extraction room and lower replacement risk). The optimum balances these effects, with the balance depending on the constraint apparatus's quality-detection capacity.

The empirical implication is that polities with strong quality-detection apparatus produce substrates with higher quality investment, even though the substrate's intrinsic incentive is to underinvest. The constraint apparatus pushes substrate toward quality not because substrate values quality intrinsically but because the constraint apparatus makes quality failure costly to substrate.

This is what the framework's prescriptions ultimately do for population welfare. They do not change the substrate's intrinsic preferences (substrate still prefers more rent and less effort); they change the constraint structure under which substrate optimizes, producing higher equilibrium quality even though substrate's preferences have not changed.

2.9 Comparative Statics

The first-order conditions produce comparative-statics results characterizing how substrate's choices respond to changes in the constraint apparatus and the polity's environment.

2.9.1 On Extraction

Higher opacity v raises optimal extraction τ^*: opacity reduces detection capacity for extraction, which reduces r_τ at any given τ, which raises optimal τ. The asymptotic limit $\tau^*(v) \to \infty$ as $v \to \infty$ obtains.

Higher coordination cost K raises optimal τ^*: harder coordination reduces r_τ at any given τ.

Higher reputation horizon ρ lowers optimal τ^*: substrate that values long-term cooperation more is more willing to forego short-term rent extraction.

Higher quality Q raises optimal τ^*. This is the new comparative-static result. Higher quality raises the revolt threshold τ_c, which expands the substrate-favorable region, which raises optimal extraction. The substantive interpretation is that quality buys extraction room.

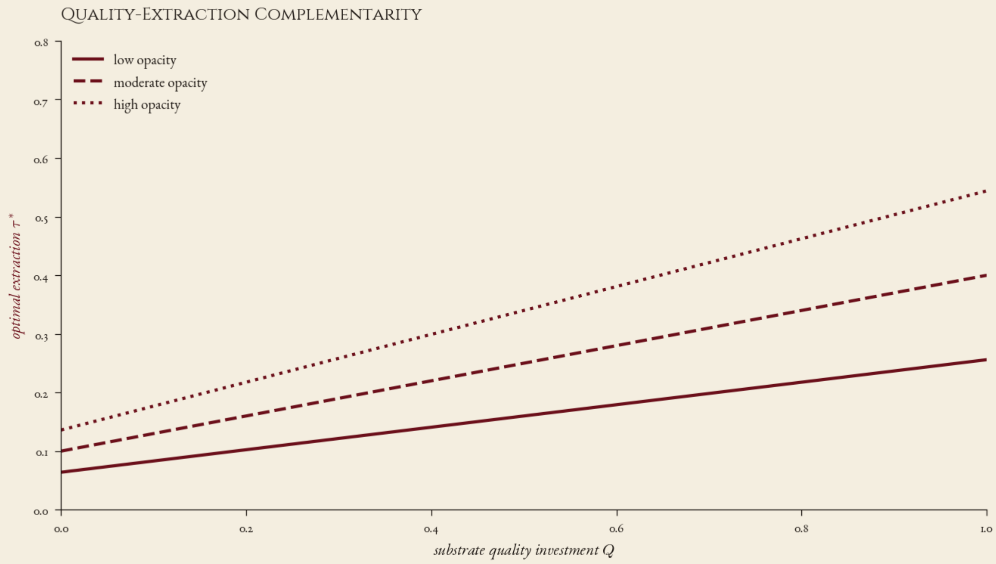

FIGURE 2.5. *Module 2: The quality-extraction complementarity result. Higher substrate quality investment Q raises optimal extraction τ^* across opacity levels. The intuition is that high-quality substrate buys higher tolerance for extraction through the population's revolt threshold.*

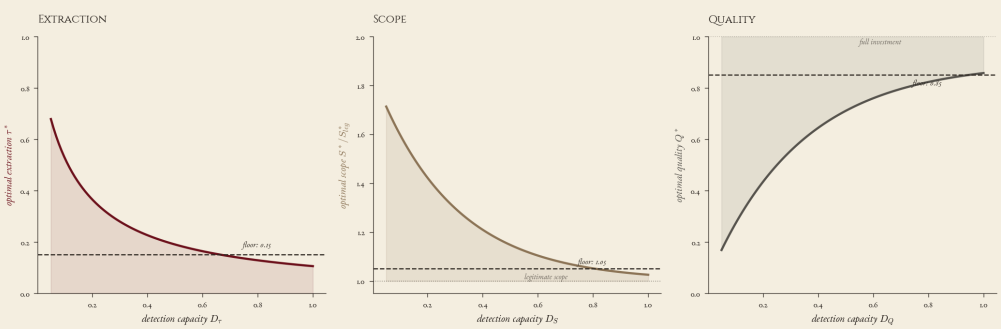

FIGURE 2.6. *Module 2: Substrate's optimal choices over extraction, scope, and quality as a function of constraint apparatus detection capacity. Extraction τ^* falls with extraction-detection capacity. Scope S^* falls with scope-detection capacity. Quality Q^* rises with quality-detection capacity. Each curve approaches a structural floor under joint implementation.*

2.9.2 On Scope

Higher constraint diversity targeting scope-detection lowers the substrate's optimal scope S^*. The constraint apparatus raises detection risk in illegitimate domains, which makes scope creep unprofitable for substrate.

Higher opacity v raises optimal scope S^*. Opacity reduces detection capacity across all dimensions, including scope detection.

Higher coordination cost K raises optimal scope S^*. Harder coordination makes population response to scope creep less effective.

Higher quality Q has ambiguous effect on optimal scope. Higher quality raises the population's tolerance for substrate authority generally, which may extend to tolerance for some scope creep. But higher quality also produces higher institutional capacity in legitimate domains, which may reduce substrate's marginal benefit from new domains. The net effect depends on whether the tolerance effect or the diminishing-returns effect dominates.

2.9.3 On Quality

Higher constraint diversity targeting quality-detection raises the substrate's optimal Q^*. The constraint apparatus raises replacement risk from quality failure, which makes quality investment more attractive.

Higher opacity v lowers optimal Q^*. Opacity reduces detection capacity for quality failure, weakening pressure on substrate to invest in quality.

Higher coordination cost K lowers optimal Q^*. Harder coordination makes population response to quality failure less effective, weakening pressure on substrate.

Higher legitimate scope S^* raises optimal Q^*. The marginal benefit of quality investment depends on the value of legitimate functions; broader legitimate scope means more value depends on quality, raising marginal benefit.

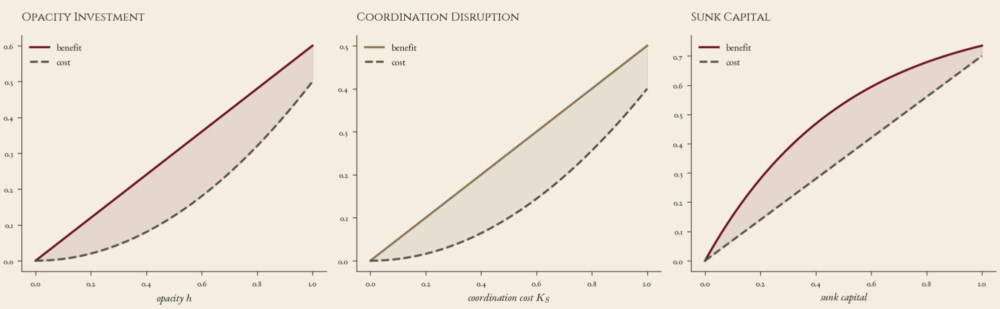

FIGURE 2.7. *Module 2: Substrate's investment choices in opacity, coordination disruption, and sunk capital. Each investment carries cost but produces favorable shifts in substrate's tolerance frontier or replacement-risk function.*

2.9.4 Cross-Dimensional Effects

The three substrate dimensions interact in the optimization in ways that produce additional substantive predictions.

Quality-extraction complementarity. As established in the extraction FOC, higher quality raises optimal extraction through the revolt threshold mechanism. This produces the empirically observed pattern of high-quality substrates extracting more (in absolute terms) than low-quality substrates. The two are linked through the tolerance frontier.

Scope-extraction substitution. Substrate has multiple ways to acquire rent. Higher τ extracts more from existing scope. Larger S creates new scope to extract from. These are partial substitutes in substrate's optimization. A polity that effectively constrains τ may produce substrate that compensates by expanding S; a polity that effectively constrains S may produce substrate that compensates by raising τ. The constraint apparatus must address both dimensions to prevent compensation.

Quality-scope reinforcement. Higher quality raises the value of legitimate functions, which raises the marginal benefit of expanding scope into adjacent domains where the existing capacity can be deployed. This produces a tendency for high-quality substrates to expand into adjacent illegitimate domains: a substrate effective at one kind of regulation has positive incentive to extend regulation to adjacent activities. The constraint apparatus must monitor scope expansion specifically to prevent this reinforcement effect.

2.10 Asymptotic Limits

The framework's apparatus produces asymptotic limits in each substrate variable. As detection capacity falls in any channel, the corresponding substrate choice moves toward its unconstrained extreme.

2.10.1 *Extraction Limit*

The asymptotic result for extraction: as opacity $v \to \infty$, optimal extraction $\tau^*(v) \to \infty$. Substrate extracts unboundedly when detection of extraction is impossible.

2.10.2 *Scope Limit*

As scope-detection capacity falls toward zero (the constraint apparatus has no capacity to detect scope creep), the substrate's optimal scope expands toward the universe of available domains:

$$\lim_{r_S \to 0} S^* = \mathcal{D}.$$

Substrate enters every domain where any rent is available, because there is no detection cost to entry. The polity loses the distinction between legitimate and illegitimate substrate scope; substrate operates everywhere.

This asymptotic limit corresponds to the empirical pattern of unconstrained-administrative-state development in polities with weak scope-detection apparatus. The administrative state expands into ever-broader domains because nothing detects or resists the expansion. The framework's prescription on resource diversity targeting scope detection is what prevents this trajectory.

2.10.3 *Quality Limit*

As quality-detection capacity falls toward zero (the constraint apparatus has no capacity to detect quality failure), the substrate's optimal quality investment falls toward zero:

$$\lim_{r_Q \to 0} Q^* = 0.$$

Substrate underinvests entirely in legitimate function performance because quality failure is undetectable and therefore unpunishable. The polity's substrate exists at appropriate scope and within extraction limits, but performs legitimate functions at zero quality.

This is the substrate dysfunction pattern characterized in Module 1 as the fourth pattern alongside the three failure modes. The asymptotic limit shows that substrate dysfunction is the equilibrium outcome when quality detection fails completely. Substrate has no intrinsic reason to invest in quality; the constraint apparatus is what produces quality investment, and its absence produces the dysfunction.

2.10.4 *Combined Limits*

The three limits operate simultaneously when the constraint apparatus fails comprehensively. As constraint apparatus capacity collapses across extraction, scope, and quality detection:

$$\tau^* \to \infty, \quad S^* \to \mathcal{D}, \quad Q^* \to 0.$$

The polity moves toward the configuration of unbounded extraction across all domains with zero quality of legitimate function performance. This is the framework's terminal-capture configuration: substrate everywhere, extracting maximally, performing nothing.

The framework's analysis of cascade dynamics in subsequent modules characterizes how a polity moves from operational equilibrium toward this terminal configuration. The asymptotic limits establish what the trajectory ends in; the cascade analysis establishes how it gets there.

2.11 The U-Shape Question

The framework's analysis includes a U-shape characterization of $\tau^*(v)$ at low opacity. At very low v, substrate hedges below the natural revolt threshold to avoid triggering coordination on revolt; at moderate v, substrate optimizes around the threshold; at high v, substrate extracts asymptotically. The result is that optimal extraction is non-monotonic in opacity: it falls as opacity rises from very low values, then rises as opacity continues to grow.

A natural question is whether analogous U-shapes obtain for scope and quality dimensions. The answer is no, and the reason is substantive: the U-shape in extraction emerges from the population's threshold-coordinated revolt dynamic, which is specific to the extraction dimension and does not generalize to scope and quality.

2.11.1 *Why Extraction Has a U-Shape*

The U-shape in extraction emerges because the population's revolt response is threshold-coordinated. At very low opacity, the population can detect substrate behavior accurately, and the revolt threshold is sharp: small increases in extraction past τ_c produce coordinated revolt. Substrate hedges below τ_c to avoid triggering coordination, and the hedging produces τ^* below the visible threshold.

At moderate opacity, the threshold becomes less sharp because population's coordination depends on signal aggregation that is impaired by noise. Substrate can optimize closer to the threshold without triggering immediate coordination, and τ^* rises.

At high opacity, the threshold disappears effectively because population cannot coordinate on substrate behavior at all. Substrate extracts asymptotically, with $\tau^* \to \infty$.

The U-shape thus reflects a specific mechanism: threshold-coordinated revolt under signal noise. The dynamic is specific to extraction because the population's response to extraction is the threshold-coordinated revolt mechanism that the framework formalizes.

2.11.2 *Why Scope Does Not Have a Comparable U-Shape*

Scope creep is typically gradual, and the population's response to scope creep is typically gradual rather than threshold-triggered. New substrate domains accumulate over time, with each individual entry being relatively small in its immediate effect but contributing to substantial cumulative scope over decades. The population does not coordinate revolt against any individual entry; the constraint apparatus detects entries through legislative debate, judicial review, and electoral consequence, with the response being gradual administrative reversal rather than threshold-triggered revolt.

Without the threshold-coordinated revolt mechanism, the U-shape dynamic does not arise. Substrate's scope choice is driven by cost-benefit analysis of each domain entry rather than by hedging against a population-level coordination threshold. The optimal scope under low scope-detection capacity is monotonically larger than under high detection capacity, with no hedging region at very low detection.

2.11.3 *Why Quality Does Not Have a Comparable U-Shape*

Quality investment is largely invisible to the population in the short run, and the population's response to quality failure is typically delayed and not threshold-triggered. Quality decisions involve internal substrate operations that the population observes only through their downstream effects on legitimate function performance, which manifest gradually rather than as a step change.

Without the threshold-coordinated revolt mechanism, the U-shape dynamic does not arise for quality either. Substrate's quality investment is driven by the constraint apparatus's detection of quality failure, which produces gradual replacement risk through electoral and other channels rather than threshold-triggered revolt. Optimal quality investment is monotonically lower under low quality-detection capacity than under high detection capacity, with no hedging dynamic.

2.11.4 *What This Means for the Framework*

The U-shape result is preserved as an extraction-specific characterization. Scope and quality have their own equilibrium dynamics, characterized by the comparative-statics and asymptotic limits derived in previous sections, but these dynamics do not exhibit non-monotonic behavior analogous to the extraction U-shape.

The framework's analytical position is therefore that substrate's behavioral variables have different response functions to constraint apparatus capacity. Extraction exhibits U-shape behavior at low opacity due to threshold-coordinated revolt. Scope and quality exhibit monotonic behavior in their respective detection capacities. The framework's predictions are richer in extraction (where the U-shape produces non-trivial hedging at low opacity) and simpler in scope and quality (where the relationships are monotonic).

This produces a substantive prediction. Polities can have meaningfully different substrate behavior across the variables even at similar overall constraint apparatus capacity. A polity with strong extraction-detection but weak scope-detection produces substrate that extracts modestly but expands into many illegitimate domains. A polity with strong scope-detection but weak quality-detection produces substrate that operates within legitimate scope but performs poorly. Extraction, scope, and quality are partially decoupled, and the constraint apparatus must address each separately rather than relying on overall constraint capacity.

2.12 Substrate Utility Versus Population Utility

The framework distinguishes substrate utility from population utility explicitly. A simpler treatment might collapse these into a single optimization with population's interests showing up only through replacement risk; the explicit distinction is required for the analysis of substrate quality and the funding-mechanism comparison.

2.12.1 *Substrate Utility*

Substrate's utility function is the substrate's discounted rent stream, with the substrate optimizing its own consumption of extracted resources subject to constraints. Substrate values rent positively: rent is what substrate keeps from extracted resources, after allocating some fraction to investments (opacity, coordination disruption, quality, sunk capital).

Substrate's utility does not directly value population welfare. Substrate values quality investment only instrumentally, through its effect on tolerance frontier and replacement risk. Substrate values legitimate scope only instrumentally, through its effect on scope creep detection. The substrate's intrinsic preferences are toward higher rent, lower investment, and lower replacement risk, with the constraint apparatus producing pressure that shifts substrate's choices toward population-favorable outcomes.

2.12.2 *Population Utility*

Population's utility depends on what substrate provides through quality investment minus what substrate extracts. Specifically:

$$U_P = Q \cdot V(S^*) \cdot \phi - \tau \cdot V(S),$$

where $Q \cdot V(S^*) \cdot \phi$ is the value population receives from legitimate function performance (with ϕ a multiplier reflecting the value of quality investment) and $\tau \cdot V(S)$ is what population pays to substrate through extraction.

Population utility depends positively on quality investment (population gets more from substrate that performs legitimate functions effectively), negatively on extraction (substrate keeps more of what population would otherwise have), and ambiguously on scope. Population utility depends on whether substrate's scope extends into legitimate domains where its services are valuable or into illegitimate domains where its services are net harms.

2.12.3 *The Constraint Apparatus's Role*

The constraint apparatus does not change substrate's utility function. Substrate continues to optimize its own discounted rent stream regardless of constraint apparatus configuration. What the constraint apparatus does is change the constraints under which substrate optimizes, which shifts the equilibrium outcome.

In the framework's analysis, the constraint apparatus's role is to align substrate's optimization with population utility. Specifically:

Extraction-detecting constraint resources raise the cost to substrate of high extraction, shifting optimal τ downward. Population benefits because it pays less to substrate.

Scope-detecting constraint resources raise the cost to substrate of scope creep, keeping optimal S closer to S^*. Population benefits because substrate operates only in domains where its services are valuable.

Quality-detecting constraint resources raise the cost to substrate of quality failure, shifting optimal Q upward. Population benefits because substrate performs legitimate functions effectively.

The constraint apparatus's effect on substrate is therefore aligned with population's interests on each behavioral variable. Joint implementation of the four prescriptions produces a constraint apparatus that addresses each variable, producing substrate behavior closer to population-optimal across the board.

2.12.4 *The Co-Evolutionary Dynamics*

The constraint apparatus is itself maintained through population virtue (Module 1). Population virtue is itself produced by substrate-equivalent institutions whose quality dimension affects what they produce. The framework's apparatus thus connects substrate behavior to constraint apparatus to population virtue to virtue-producing institutions, with all four layers interacting.

Population's effective utility depends on substrate behavior, which depends on constraint apparatus capacity, which depends on population virtue, which depends on virtue-producing institutions, which are themselves substrate-equivalent and subject to substrate-quality dynamics.

This produces a multi-layered analytical structure that the framework's prescriptive program addresses. The four prescriptions operate primarily on the constraint apparatus layer. Their effect propagates through the substrate behavior layer (lower extraction, less scope creep, higher quality) and through the population utility layer (more from substrate, less paid to substrate). The prescriptions also operate indirectly on the virtue-producing institutions layer through their effect on the constraint apparatus that protects those institutions.

2.12.5 *The Funding Question Briefly Addressed*

The cost-bearing structure of substrate's quality investment varies across institutional configurations. The full comparative analysis is developed in Chapter 5. Here, three points:

First, in standard substrate configurations, the population bears the cost of substrate's quality investment through the underlying extraction. The substrate is allocating extracted resources, not bearing costs of its own.

Second, in left-anarchist configurations, the cost is borne by population through voluntary contributions to substrate-equivalent institutions. The funding mechanism differs but the cost-bearing structure is similar.

Third, in anarcho-capitalist configurations, the cost is borne by users of substrate-equivalent services through market fees. The framework's deeper substrate-smuggling diagnosis applies particularly to virtue-producing institutions in this configuration: their costs are not addressed in the position's account because the institutions are not acknowledged as substrate-equivalent.

The substantive comparison across configurations belongs in Module 7. The point for Module 2 is that substrate's optimization is over allocation of population-extracted resources, with the underlying cost borne by the population regardless of configuration.

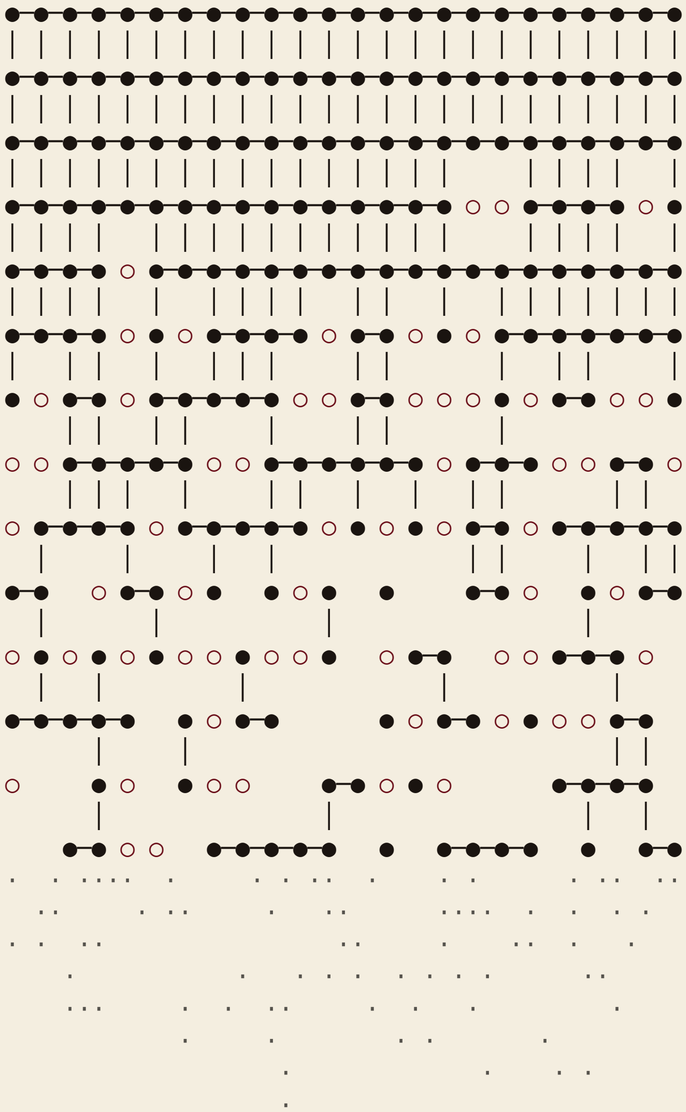

CHAPTER III

Legitimacy Resources and Constraint Capture

This chapter develops the dynamics of legitimacy resources and the topology of the constraint apparatus. Module 3 treats legitimacy resources as common-pool stocks subject to substrate capture, with the cascading depletion theorem and the uneven cascade pattern. Module 4 develops the coverage graph topology, multi-resource detection theorem, and cascade vulnerability characterization.

The chapter's contributions are several. The legitimacy resource portfolio is characterized in terms of which detection functions each resource supports, with resource-to-parameter mappings reflecting these roles. The cascading depletion theorem characterizes cascade across substrate dimensions, with the substantive observation that cascade can proceed unevenly formalized as Proposition 3.5.1. The coverage graph carries edge modality, producing subgraphs G_τ, G_S, G_Q that the framework analyzes separately and together. The multi-resource detection theorem applies dimension by dimension, with conditions for detection requiring paths through resources that evaluate the relevant dimension. The cascade vulnerability theorem is dimension-specific, allowing characterization of polities where cascade vulnerability differs across dimensions.

Part One: Module 3 (Legitimacy Resources)

3.1 Legitimacy Resources

Module 1 established that substrate is necessary for voluntary cooperation at scale and identified population virtue as a substantive consideration affecting substrate-population dynamics. Module 2 established substrate as a strategic actor optimizing across extraction, scope, and quality dimensions, with the constraint apparatus producing pressure that aligns substrate's optimization with population interests. This module characterizes the resources that constitute the constraint apparatus and analyzes their dynamics over operational time.

3.1.1 The Resource Portfolio

Legitimacy resources are the institutional and informational stocks that support the population's capacity to evaluate and constrain substrate behavior. The framework identifies six resource categories:

Founding mythology. The shared narrative about why substrate exists and what its legitimate purposes are. Founding mythology provides the substantive standard against which substrate behavior is evaluated. A polity with strong founding mythology has clear criteria for distinguishing legitimate substrate activity from substrate overreach.

Procedural authority. The institutional infrastructure for evaluating substrate behavior through formal processes: judicial review, legislative oversight, administrative procedure, electoral mechanisms. Procedural authority provides the structured channels through which substrate evaluations are conducted.

Expertise authority. The professional and technical communities that evaluate substrate behavior in their respective domains: medical professional bodies, engineering societies, scientific academies, legal professional organizations, economic and policy research institutions. Expertise authority provides the substantive evaluative capacity in domains where general public assessment is inadequate.

Democratic authority. The institutional and cultural infrastructure for population participation in substrate evaluation: voting mechanisms, civil society organizations, deliberative forums, civic education. Democratic authority provides the population-aggregation capacity for substrate evaluation.

Traditional authority. The historical-cultural patterns that articulate what cooperation looks like and what substrate behavior is appropriate: religious institutions, cultural traditions, family structures, community practices. Traditional authority provides the deep value substrate against which immediate substrate behavior is evaluated.

Crisis authority. The institutional capacity to conduct extraordinary evaluation of substrate behavior when normal procedures are inadequate: investigative journalism in moments of crisis, special prosecutors, congressional hearings on specific events, constitutional convention mechanisms. Crisis authority provides the surge capacity for substrate evaluation when standard mechanisms fail.

Let $\mathbf{L}_t = (L_{1,t}, L_{2,t}, \ldots, L_{6,t}) \in \mathbb{R}^6_+$ denote the vector of legitimacy resource stocks at period t.

Each component $L_{j,t}$ measures the operational capacity of resource j, with higher values indicating more functional resources.

3.1.2 Each Resource's Detection Roles

The framework characterizes each legitimacy resource in terms of which detection functions it supports.

Founding mythology supports primarily scope detection. The shared narrative about substrate's legitimate purposes provides the criteria for distinguishing legitimate from illegitimate substrate scope. Founding mythology contributes secondarily to extraction detection (substrate's legitimate purposes typically include implicit limits on extraction) and to quality detection (legitimate purposes presuppose adequate performance).

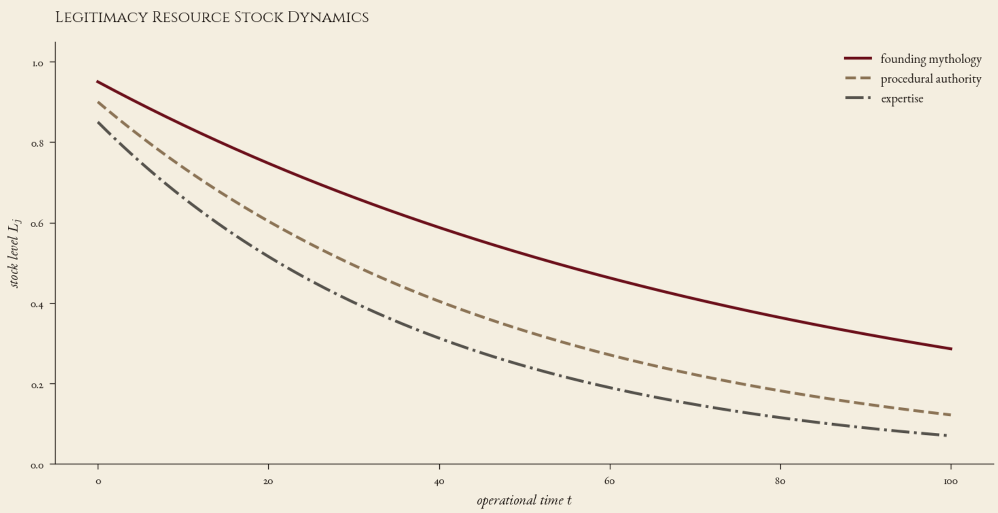

FIGURE 3.1. *Module 3: Legitimacy resource stock dynamics. Each resource has its own depletion through capture and refresh through population effort, with cross-resource dynamics through the cascading depletion mechanism.*

Procedural authority supports detection in extraction, scope, and quality. Judicial review evaluates scope (constitutional limits on substrate authority), extraction (statutory limits on what substrate can demand), and quality (adequacy of substrate's legitimate function performance). Legislative oversight similarly addresses each behavioral channel through committee review, budget appropriation, and confirmation processes. Procedural authority is the most multifunctional of the legitimacy resources.

Expertise authority supports primarily quality detection. Professional bodies evaluate whether substrate is performing its functions effectively in their domains. Medical professional bodies evaluate health policy quality. Engineering societies evaluate infrastructure policy quality. Scientific academies evaluate research and development policy quality. Expertise authority contributes secondarily to scope detection (professional bodies have views on what counts as legitimate substrate scope in their domains) and to extraction detection (professional standards include views on appropriate compensation and resource use).

Democratic authority supports detection in all three behavioral channels through population aggregation. Voting mechanisms allow population to express judgments about substrate performance on extraction, scope, or quality, with electoral consequence operating as a general-purpose evaluation channel. Civil society organizations aggregate population concerns about specific substrate behavior on each variable.

Traditional authority supports primarily scope detection through articulating what cooperation looks like and what substrate behavior is culturally appropriate. Religious institutions articulate moral limits on substrate scope. Cultural traditions articulate what kinds of substrate activity are acceptable in the polity's specific context. Traditional authority contributes secondarily to quality detection (traditional standards for proper performance) and to extraction detection (traditional norms about appropriate substrate resource use).

Crisis authority supports detection in any of the three behavioral channels in extraordinary circumstances. Investigative journalism, special prosecutors, congressional hearings, and constitutional convention mechanisms can evaluate substrate behavior across the full spectrum when standard mechanisms fail. Crisis authority is the surge capacity that allows the constraint apparatus to address substrate behavior that has escaped routine evaluation.

3.1.3 Internal-Substrate Versus External Constraint Resources

The six legitimacy resource categories cut across a distinction the framework will use throughout subsequent analysis: between constraint resources internal to substrate and constraint resources external to substrate. Internal-substrate constraint resources are institutions whose function is to constrain substrate behavior but which are themselves part of substrate: judicial review of legislative and executive action (the judiciary constrains other parts of substrate), professional civil service rules insulating bureaucracy from political pressure, internal audit and inspector-general functions, central bank independence from political direction, statistical agency independence, prosecutorial independence within the executive. External constraint resources are institutions outside substrate that evaluate substrate behavior: independent press, civic society organizations, professional bodies with no formal substrate role, religious institutions evaluating substrate scope, academic institutions producing analysis substrate cannot direct.

Each of the six resource categories above contains both internal and external instances. Procedural authority contains internal-substrate procedures (judicial review of substrate action) and external procedures (citizen petition processes, ballot initiative mechanisms). Expertise authority contains internal-substrate expertise (statistical agencies, central bank professional staff) and external expertise (independent professional associations, academic disciplines). Democratic authority contains internal mechanisms (electoral commissions within substrate) and external mechanisms (party organizations, civic mobilization).

The distinction matters analytically because internal-substrate constraint resources are subject to a failure mode the framework will examine: substrate has standing capacity to degrade them through internal politicization, suppression of dissent within substrate, and sycophantic feedback structures that prevent accurate self-perception. External constraint resources can be captured, but capture requires substrate action against an institution it does not formally control, which is a different mechanism than internal degradation. The cascading depletion theorem (Section 3.4) applies to both internal and external resources, but the specific dynamics differ. Subsequent sections note where the internal/external distinction modifies the analysis.

3.1.4 The Detection Capacity Vector

The constraint apparatus's detection capacity therefore has three components, each supported by different combinations of legitimacy resources. Define the detection capacity vector $\mathbf{D}(\mathbf{L}) = (D_\tau(\mathbf{L}), D_S(\mathbf{L}), D_Q(\mathbf{L}))$, where D_τ is detection capacity for extraction, D_S for scope creep, and D_Q for quality failure. Each component depends on the legitimacy resource stocks through the dimensional roles described above:

$$D_\tau(\mathbf{L}) = f_\tau(L_1, L_2, L_3, L_4, L_5, L_6) \tag{3.1}$$

$$D_S(\mathbf{L}) = f_S(L_1, L_2, L_3, L_4, L_5, L_6) \tag{3.2}$$

$$D_Q(\mathbf{L}) = f_Q(L_1, L_2, L_3, L_4, L_5, L_6) \tag{3.3}$$

The functions f_τ, f_S, f_Q are increasing in each L_j but with weights reflecting which dimension the resource primarily supports. Founding mythology and traditional authority weight scope detection most heavily. Expertise authority weights quality detection most heavily. Procedural and democratic authority contribute roughly equally to all three components. Crisis authority contributes to all three through surge capacity rather than continuous evaluation.

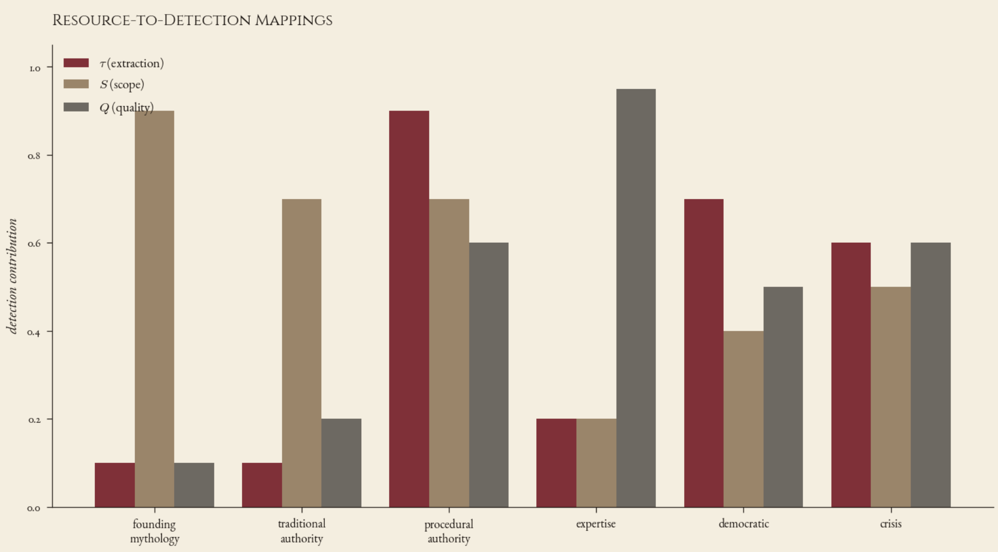

FIGURE 3.2. *Module 3: Resource-to-parameter mappings. Each legitimacy resource contributes to detection capacity in one or more behavioral channels, with the contribution magnitude characterized through the resource's primary and secondary detection roles.*

The substantive observation is that detection capacity in any dimension depends on multiple legitimacy resources, with no single resource being sufficient for any dimension's detection. This is what makes the resource portfolio a portfolio rather than a single resource: the dimensions of detection are sustained by combinations of resources operating in different modalities.

3.2 Resource-to-Parameter Mappings

The substrate optimization developed in Chapter 2 depends on parameters that characterize the constraint environment. Resource stocks map to extraction-relevant parameters (opacity v, coordination cost K, and reputation horizon ρ) and to scope and quality parameters.

3.2.1 *Extraction-Relevant Mappings*

The mappings to extraction-relevant parameters reflect which legitimacy resources support extraction detection.

Opacity $v(\mathbf{L})$ is decreasing in resources that support extraction detection: procedural authority (especially press freedom and judicial review), expertise authority (in domains where extraction is technical), democratic authority (electoral monitoring), and crisis authority (investigative surge capacity). Founding mythology and traditional authority contribute less to extraction detection and therefore have weaker effect on v.

Coordination cost $K(\mathbf{L})$ is decreasing in resources that support population coordination: democratic authority directly, civil society organization through democratic and traditional authority, communication infrastructure through procedural authority. Founding mythology and expertise authority contribute less directly.

Reputation horizon $\rho(\mathbf{L})$ is increasing in resources that support institutional memory and forward-looking commitment: founding mythology, traditional authority, and certain forms of expertise authority that maintain longitudinal records.

3.2.2 *Scope-Relevant Mappings*

Scope detection capacity depends on resources that support evaluation of whether substrate is operating within legitimate scope.

Scope detection probability $\pi_S(\mathbf{L})$ is increasing in resources that articulate legitimate scope and detect deviations: founding mythology (substantive criteria for legitimate scope), traditional authority (cultural-historical patterns of legitimate substrate activity), procedural authority (constitutional and statutory limits), and crisis authority (extraordinary review of scope creep). Expertise authority contributes secondarily through professional views on legitimate scope in specific domains.

Scope reversal capacity $\eta_S(\mathbf{L})$ is increasing in resources that support reversing detected scope creep: procedural authority (judicial nullification, legislative repeal), democratic authority (electoral consequence for scope expansion), and crisis authority (extraordinary action). The asymmetry between entry and exit costs (Module 2) makes reversal harder than detection, with reversal capacity depending more heavily on procedural authority than detection capacity does.

3.2.3 *Quality-Relevant Mappings*

Quality detection capacity depends on resources that support evaluation of whether substrate is performing legitimate functions effectively.

Quality detection probability $\pi_Q(\mathbf{L})$ is increasing in resources that evaluate legitimate function performance: expertise authority primarily (professional bodies in each domain), procedural authority (administrative review, audit functions), democratic authority (electoral consequence for poor performance), and crisis authority (investigative review of failures). Founding mythology and traditional authority contribute secondarily through standards for what adequate performance looks like.

Quality response capacity $\eta_Q(\mathbf{L})$ is increasing in resources that support correcting detected quality failure: procedural authority (administrative reform, legislative restructuring), expertise authority (professional standards enforcement), democratic authority (electoral pressure for improvement), and crisis authority (extraordinary intervention).

3.2.4 *The Resulting Parameter Vector*

The substrate's optimization parameters now form a vector reflecting the constraint apparatus's capacities for each behavioral channel:

$$\boldsymbol{\Theta}(\mathbf{L}) = (v(\mathbf{L}), K(\mathbf{L}), \rho(\mathbf{L}), \pi_S(\mathbf{L}), \eta_S(\mathbf{L}), \pi_Q(\mathbf{L}), \eta_Q(\mathbf{L})).$$

The substrate's optimal choices τ^*, S^*, Q^* from Module 2 depend on this parameter vector. As legitimacy resources deplete, the parameter vector shifts in the substrate-favorable direction for each variable.

3.3 The Substrate's Portfolio Capture Optimization

The substrate has the capacity to capture legitimacy resources, biasing each resource's evaluation function toward substrate-favorable outcomes. The substrate's portfolio capture optimization is over capture rates for each resource, with the optimization considering current rent gains, future rent capacity, and dynamic interactions among resources.

3.3.1 *The Capture Choice Variables*

Let $c_{j,t} \in [0, 1]$ denote the substrate's capture rate for resource j at period t. Capture is a continuous variable rather than binary: substrate biases evaluation functions to varying degrees rather than fully replacing them. The capture function for resource j is

$$L_{j,t+1} = (1 - c_{j,t}) \cdot L_{j,t} + R_j(\mathbf{L}_t, e_{j,t}),$$

where R_j is the population's refresh capacity for resource j given the resource portfolio and population effort $e_{j,t}$. The capture rate $c_{j,t}$ depletes the resource; the refresh function R_j replenishes it.

3.3.2 *The Substrate's Capture Optimization*

The substrate's capture decisions are part of the broader optimization developed in Module 2. Capturing resource j benefits substrate through the resulting parameter shifts: lower opacity-detection, lower scope-detection, lower quality-detection, depending on which dimensions resource j supports.

The first-order condition for substrate's capture rate of resource j is

$$\frac{\partial \Pi_S}{\partial c_j} - \delta_S \cdot \frac{\partial V_S(\mathbf{L}')}{\partial L_j'} \cdot \frac{\partial L_j'}{\partial c_j} = 0.$$

The first term is the marginal current-period benefit from capturing resource j, which depends on which substrate optimization dimensions resource j affects. The second term is the marginal future cost of depleting resource j, weighted by the substrate's discount factor.

The current-period benefit decomposes through the resource's dimensional contributions:

$$\frac{\partial \Pi_S}{\partial c_j} = \frac{\partial \Pi_S}{\partial \tau^*} \cdot \frac{\partial \tau^*}{\partial L_j} \cdot \frac{\partial L_j}{\partial c_j} + \frac{\partial \Pi_S}{\partial S^*} \cdot \frac{\partial S^*}{\partial L_j} \cdot \frac{\partial L_j}{\partial c_j} + \frac{\partial \Pi_S}{\partial Q^*} \cdot \frac{\partial Q^*}{\partial L_j} \cdot \frac{\partial L_j}{\partial c_j}.$$

Each term reflects how capture of resource j affects one of substrate's optimal choices, with the magnitude depending on which detection functions j supports. Capturing a resource that primarily supports scope detection (e.g., founding mythology) primarily affects S^*. Capturing one that primarily supports quality detection (e.g., expertise authority) primarily affects Q^*. Capturing one that supports detection in all three channels (e.g., procedural authority) affects each.

3.4 The Cascading Depletion Theorem

Depletion of any one legitimacy resource accelerates depletion of others through three mechanisms: portfolio re-optimization, refresh efficiency reduction, and complementarity reinforcement. The cascade affects all three substrate dimensions.

Theorem 3.1 *(Cascading Depletion Across Dimensions). As any resource j depletes, the substrate's optimal capture rate of remaining resources $k \neq j$ increases. Formally, in the relevant range:*

$$\frac{\partial c_k^*}{\partial L_j} < 0 \quad \textit{for } k \neq j.$$

The depletion of one resource accelerates depletion of others. The cascade dynamics affect detection capacity in each behavioral channel, with each channel's detection capacity declining as the resource portfolio depletes.

Proof. The substrate's first-order condition for capture rate c_k is

$$F_k(\mathbf{c}, \mathbf{L}) \equiv \frac{\partial \Pi_S}{\partial c_k}(\mathbf{c}, \mathbf{L}) - \delta_S \cdot \frac{\partial V_S}{\partial L_k'}(\mathbf{L}'(\mathbf{c}, \mathbf{L})) \cdot \frac{\partial L_k'}{\partial c_k}(\mathbf{c}, \mathbf{L}) = 0.$$

At an interior optimum $c_k^*(\mathbf{L})$, $F_k(\mathbf{c}^*(\mathbf{L}), \mathbf{L}) = 0$ for all k. Differentiating implicitly with respect to L_j for $j \neq k$ and applying the implicit function theorem:

$$\sum_m \frac{\partial F_k}{\partial c_m} \cdot \frac{\partial c_m^*}{\partial L_j} + \frac{\partial F_k}{\partial L_j} = 0.$$

For interior optima with strictly concave substrate objective, the matrix $[\partial F_k / \partial c_m]$ is negative definite at the optimum, with negative diagonal entries. Treating cross-effects as small relative to own-effects in the substantive range:

$$\frac{\partial c_k^*}{\partial L_j} \approx -\frac{\partial F_k / \partial L_j}{\partial F_k / \partial c_k}.$$

The denominator is negative. The sign of $\partial c_k^* / \partial L_j$ is therefore the sign of $\partial F_k / \partial L_j$. We show this is positive, which establishes $\partial c_k^* / \partial L_j < 0$.

Three components contribute to $\partial F_k / \partial L_j$, and each operates across multiple substrate dimensions.

Portfolio re-optimization channel. The current-period benefit of capturing k depends on substrate's optimal choices τ^*, S^*, Q^*, which depend on the parameter vector $\mathbf{\Theta}(\mathbf{L})$ through the resource-to-parameter mappings. Increasing L_j shifts $\mathbf{\Theta}$ in the direction that lowers substrate's optimal choices on whichever channels resource j supports. The marginal benefit of capturing k depends on the gain available, which falls when complementary resources' contributions are intact. Therefore $\partial^2 \Pi_S / \partial c_k \partial L_j > 0$ for whichever substrate variable k supports: depletion of j raises the marginal benefit of capturing k.

Refresh-cost channel. The refresh function R_j depends on coordination capacity, which depends on legitimacy resources. When L_j falls (where j is a coordination-supporting resource), refresh efficiency μ_k for other resources falls. The substrate's effective marginal capture cost decreases, which raises $\partial F_k / \partial L_j$.

Continuation-value channel. The continuation derivative $\partial V_S / \partial L_k'$ depends on $\mathbf{L}'$ through the dynamic programming recursion. When complementary resources are depleted, the marginal cost of further capture (in terms of foregone future capture opportunities) is lower. This contributes positively to $\partial F_k / \partial L_j$.

Summing the three contributions across the substrate variables that k supports, $\partial F_k / \partial L_j > 0$. Therefore $\partial c_k^* / \partial L_j < 0$ for $k \neq j$, completing the proof of the basic cascade result.

The corollary on detection capacity follows directly. The detection capacity vector $\mathbf{D}(\mathbf{L})$ is monotonically increasing in each L_j. As resources deplete, $\mathbf{D}$ falls in each component, with the rate of fall depending on which resources are depleting and which components those resources support. Cascade dynamics therefore affect each behavioral channel, with the specific pattern depending on the order and rate of depletion across resources. ■

3.5 Uneven Cascade Across Dimensions

The cascading depletion theorem establishes that cascade affects each detection channel, but the rate at which each channel declines depends on which resources are depleting. This produces a substantively important pattern: cascade can proceed unevenly, with one channel's detection capacity collapsing rapidly while others remain partially functional.

PROPOSITION 3.2 (Uneven Cascade Pattern). *The rate of detection capacity decline in dimension* d *depends on the depletion rates of resources whose primary contribution is to dimension* d*. If a polity's cascade dynamics deplete resources whose primary contribution is to one dimension faster than resources whose primary contribution is to other dimensions, the dimension whose primary supporting resources deplete fastest experiences the most rapid detection capacity decline.*

The substantive consequence is that polities can experience meaningfully different cascade patterns. A polity in which substrate first captures expertise authority (e.g., through influence over professional bodies, suppression of dissent, control of credentialing) experiences rapid quality-detection decline while scope-detection and extraction-detection remain functional initially. A polity in which substrate first captures procedural authority (e.g., through judicial appointments, legislative coalition management, administrative procedure manipulation) experiences rapid scope-detection decline. A polity in which substrate first captures democratic authority (e.g., through electoral manipulation, civil society co-optation) experiences general decline across dimensions but with surge-capacity through crisis authority potentially preserved.

3.5.1 Historical Patterns of Uneven Cascade

The historical record exhibits patterns matching this prediction.

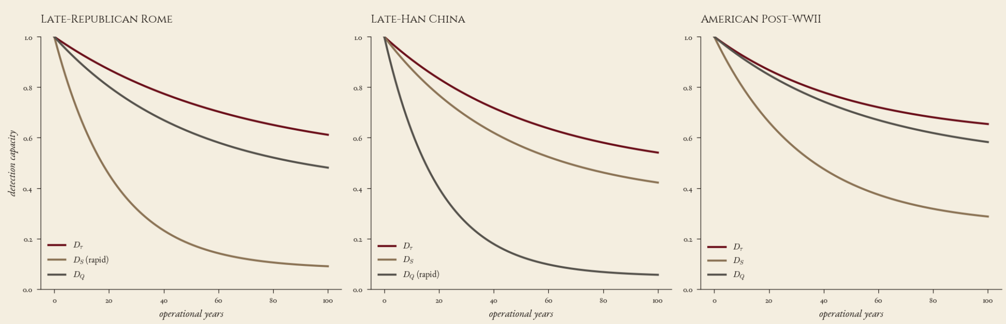

FIGURE 3.3. *Module 3: Three historical patterns of uneven cascade across dimensions. Late-Republican Rome exhibited rapid cascade in scope-detection while extraction and quality detection persisted longer. Late-Han China exhibited rapid cascade in quality-detection. The American post-WWII trajectory shows a mixed cascade pattern. Each polity's cascade is characterized by which dimensions' supporting resources captured fastest.*

Late-Republican Rome experienced rapid cascade in scope-creep detection. The substrate's expansion into imperial administration was largely undetected and unresisted as the constraint apparatus failed to maintain criteria for legitimate scope. Extraction detection within Rome itself remained more constrained for longer because the procedural authority associated with senatorial review continued to function for extraction-related matters even as it failed for scope expansion. The cascade was uneven in the predicted direction: founding mythology and traditional authority (which primarily supported scope detection) depleted faster than procedural authority's extraction-related functions.

Late-Han China experienced rapid cascade in quality-failure detection. The bureaucracy's performance failures became unaddressable as expertise authority captured by competing factions failed to maintain professional standards. Formal extraction structures continued longer because the procedural authority for tax administration remained partially functional. The cascade was uneven in the predicted direction: expertise authority (which primarily supported quality detection) depleted faster than procedural authority's extraction-related functions.

The American post-World-War-II trajectory exhibits a more recent pattern. The constraint apparatus's scope-detection capacity has depleted substantially over the period, with substrate expansion into domains beyond traditional rights-protective functions occurring with decreasing constitutional and legislative resistance. Extraction-detection has remained relatively more functional through electoral mechanisms and budget oversight, though with increasing capture. Quality-detection has been variable across domains, with strong professional bodies in some domains (medical, engineering) and weak professional bodies in others (educational policy, regulatory science). The cascade pattern is uneven across dimensions, with founding mythology and traditional authority depleting faster than expertise authority in some domains.

3.5.2 Implications for Empirical Analysis

The uneven cascade prediction produces falsifiable empirical claims. A polity's specific cascade trajectory can be characterized by tracking which legitimacy resources deplete at what rates and by examining whether the resulting detection capacity decline matches the predicted dimensional pattern.

The uneven cascade also produces a substantive observation about the framework's prescriptions. Joint implementation of the four prescriptions (Module 7) is required precisely because uneven cascade can produce dimension-specific failures even when overall constraint capacity appears adequate. A polity that maintains strong extraction-detection apparatus may still produce poor outcomes if its scope-detection or quality-detection apparatus has cascaded. The framework's apparatus identifies the dimension-specific failures that aggregate measures of constraint capacity miss.

3.6 The Population's Refresh Problem

The population's refresh problem is a collective action problem: each agent's individual incentive to contribute to refresh effort produces equilibrium refresh below the social optimum. The framework's analysis characterizes how refresh effort is allocated across resources supporting different detection channels and what this implies for the prescriptive program.

3.6.1 Refresh as Multi-Dimensional Effort Allocation

Population refresh effort is allocated across the six legitimacy resources, with each unit of effort contributing to specific resources. The population's refresh problem is therefore not just a single collective action problem but a multi-dimensional allocation problem: how to distribute refresh effort across resources to maintain detection capacity across dimensions.

The standard collective action analysis applies to refresh of any individual resource. Each agent's marginal benefit of contributing to refresh of resource j is the per-agent share of the value produced by resource j. The marginal cost is the agent's effort cost. Equilibrium refresh of resource j occurs where these balance, with the equilibrium contribution generally falling below the social optimum due to the standard free-rider gap.

3.6.2 *The Allocation Problem Across Resources*

Within the framework's apparatus, population's refresh allocation across the six resources is itself subject to coordination problems. Resources whose value is more visible to population (procedural authority through judicial decisions, democratic authority through electoral mechanisms) tend to receive more refresh effort than resources whose value is less visible (founding mythology, traditional authority operating through cultural transmission). The allocation pattern produces uneven refresh that affects which dimensions of detection capacity are maintained.

The substantive observation is that refresh effort tends to flow toward resources with visible value, leaving resources with less visible value relatively under-refreshed. This produces a pattern in which scope-detection and quality-detection (which depend more on founding mythology, traditional authority, and expertise authority that operates through less visible mechanisms) tend to be under-refreshed relative to extraction-detection (which depends more on procedural and democratic authority operating through visible mechanisms).

The historical pattern of substrate behavior matches this prediction. In polities where refresh effort concentrates on visible mechanisms, the substrate experiences less constraint on scope and quality dimensions over operational time. The framework's third prescription (refresh capacity) is the institutional response: design refresh mechanisms specifically for the less-visible resources whose refresh would otherwise lag.

3.6.3 *Ostrom's Design Principles Across Dimensions*

Ostrom's design principles for sustainable commons translate to the legitimacy commons across all detection channels, but the specific institutional features required differ by channel.

For extraction detection, Ostrom's principles support visible monitoring (transparency in budgets, public reporting of substrate revenue), graduated sanctions (electoral consequence for excessive extraction, judicial review for unauthorized extraction), conflict-resolution forums for disputes about appropriate extraction levels.

For scope detection, Ostrom's principles support clear boundary specification (constitutional articulation of legitimate scope, founding mythology as substantive criteria), monitoring through institutional memory (traditional authority, historical records), and graduated response to scope creep (judicial nullification, legislative repeal).

For quality detection, Ostrom's principles support expertise communities with formal continuity (professional bodies, scientific academies), peer evaluation mechanisms (professional standards enforcement, accreditation), and coordination forums for cross-professional evaluation when scope and quality issues intersect.

The framework's prescriptive program (Chapter 5) operationalizes these design principles across each detection channel, producing institutional recommendations that address refresh capacity for each separately.

Part Two: Module 4 (Coverage Graph Topology with Dimensional Structure)

3.7 Constraint Resources and Dimensional Evaluation Functions

Module 3 characterized legitimacy resources broadly, with each resource supporting detection in one or more substrate dimensions. This module isolates constraint resources within the broader portfolio. Constraint resources are those whose function is specifically to evaluate substrate behavior, as opposed to legitimacy resources that perform other functions (founding mythology articulating purposes, traditional authority transmitting values).

3.7.1 The Set of Constraint Resources

Let $\mathcal{C} \subseteq \{1, 2, \ldots, J\}$ denote the set of constraint resources within the legitimacy resource portfolio. Constraint resources include:

Procedural authority resources: federal courts (district, appellate, supreme), state courts, specialized courts (tax, foreign intelligence surveillance, international trade), legislative committees (House and Senate, with their subcommittees), inspector general offices, the Government Accountability Office, the Congressional Budget Office, administrative law judges.

Independent press resources: major newspapers, magazines of record, broadcast journalism, specialty press, wire services, regional papers, and (increasingly) independent online publications.

Professional bodies: medical professional organizations (AMA, specialty boards), legal professional organizations (ABA, state bars), engineering societies, scientific academies (the National Academies of Sciences, Engineering, and Medicine), economic and policy research institutions, accreditation bodies.

Civil society organizations: advocacy organizations across the political spectrum, watchdog organizations, legal advocacy groups, think tanks operating with formal independence from substrate.

Academic institutions: research universities, research institutes, the disciplinary structure of academic inquiry.

International institutions and observers: treaty bodies, international press, foreign governments, multilateral organizations, comparative scholarly work.

Each constraint resource is associated with an evaluation function specifying what kinds of substrate actions it evaluates and what verdict it issues.

3.7.2 *Evaluation Functions*

The framework characterizes each constraint resource's evaluation function in terms of which substrate dimensions it evaluates. Define the evaluation function for resource $k \in \mathcal{C}$ as $f_k : \mathcal{A} \times \{\tau, S, Q\} \to \mathcal{V}$, where $\mathcal{A}$ is the space of substrate actions, $\{\tau, S, Q\}$ is the set of substrate dimensions, and $\mathcal{V} = \{0, 1\}$ is the verdict space (1 = illegitimate, 0 = legitimate).

For an action $a \in \mathcal{A}$ that constitutes substrate behavior in dimension $d \in \{\tau, S, Q\}$, the evaluation function $f_k(a, d)$ is the verdict that resource k would issue regarding a's legitimacy on dimension d. The function is defined only for those combinations where k is positioned to evaluate dimension d; otherwise the function is undefined or returns a null verdict.

Federal courts evaluate scope (constitutional limits on substrate authority through Article III review), extraction (constitutional and statutory limits on what substrate can demand), and quality (whether substrate's legitimate function performance meets statutory standards). Federal courts thus participate in each channel, with different specific evaluation processes for each.

Press resources evaluate primarily extraction (transparency about substrate revenue and spending) and scope (reporting on substrate expansion into new domains), with quality evaluation being domain-specific (good investigative journalism on healthcare quality, legal system performance, regulatory effectiveness).

Professional bodies evaluate primarily quality (technical evaluation of substrate's domain performance), with secondary evaluation of scope (professional views on what counts as legitimate substrate scope in their domains) and extraction (professional standards for compensation and resource use).

Civil society organizations have varied evaluation functions depending on the organization's specific focus. Some evaluate scope (constitutional advocacy organizations), some evaluate extraction (taxpayer advocacy organizations), some evaluate quality (policy effectiveness organizations).

Academic institutions evaluate primarily quality (research evaluation of policy effectiveness) and scope (constitutional and historical scholarship on legitimate substrate scope), with secondary evaluation of extraction.

International observers evaluate each behavioral channel through external comparison, treating the polity's substrate behavior against the patterns of comparable polities.

Each constraint resource has a specific profile of evaluation functions, with some resources contributing to detection in all three behavioral channels and others specializing in one or two.

3.7.3 *Capture and Evaluation Bias*

Capture of constraint resource k biases its evaluation functions across the dimensions where k evaluates. The capture transformation $\Phi_k : f_k \mapsto f'_k$ produces a captured evaluation function where, for actions a in the substrate-favorable region $\mathcal{S}$ and any dimension d that k evaluates:

$$f'_k(a, d) = 0 \quad \text{(legitimate)} \quad \text{for } a \in \mathcal{S},$$

even when $f_k(a, d) = 1$ in the uncaptured state. The captured resource issues legitimate verdicts in substrate-favorable cases regardless of which dimension is being evaluated.

This is the substrate capture transformation operating on the multi-dimensional evaluation function. A captured federal court fails to evaluate scope, extraction, and quality faithfully across substrate-favorable cases. A captured professional body fails to evaluate quality faithfully across substrate-favorable cases. The capture's specific effects depend on which dimensions the resource was positioned to evaluate, but the bias direction is uniformly toward substrate-favorable verdicts in the substrate-favorable region.

3.7.4 The Substrate-Favorable Region

The substrate-favorable region $\mathcal{S} \subseteq \mathcal{A}$ encompasses substrate behavior on each of its variables. Substrate prefers to be evaluated as legitimate when it is extracting heavily, expanding scope, and underinvesting in quality. The capture's substrate-favorable bias operates simultaneously on whichever channels the captured resource evaluates.

ASSUMPTION 3.3 *(Capture Inclusion Across Dimensions).* For each substrate dimension $d \in \{\tau, S, Q\}$ and each captured resource k, the substrate-favorable region $\mathcal{S}$ includes the actions a corresponding to substrate behavior that k would evaluate as illegitimate on dimension d in the uncaptured state. Capture of k therefore biases evaluation toward substrate-favorable verdicts on dimension d for all actions a that would otherwise be detected.

This assumption says that capture is uniformly substrate-favorable across whatever dimensions the captured resource was positioned to evaluate. The assumption is substantively defensible: substrate captures resources to bias their verdicts in substrate's favor across all the dimensions where the resource matters.

3.8 The Coverage Graph with Dimensional Structure

3.8.1 Coverage Graph Definition

The coverage graph $G_{\mathcal{C}}$ is a directed graph with vertex set $\mathcal{C}$ (constraint resources) and edge set $E_{\mathcal{C}}$. An edge runs from k to j if k is positioned to evaluate j's evaluations or to evaluate substrate actions that j would also evaluate.

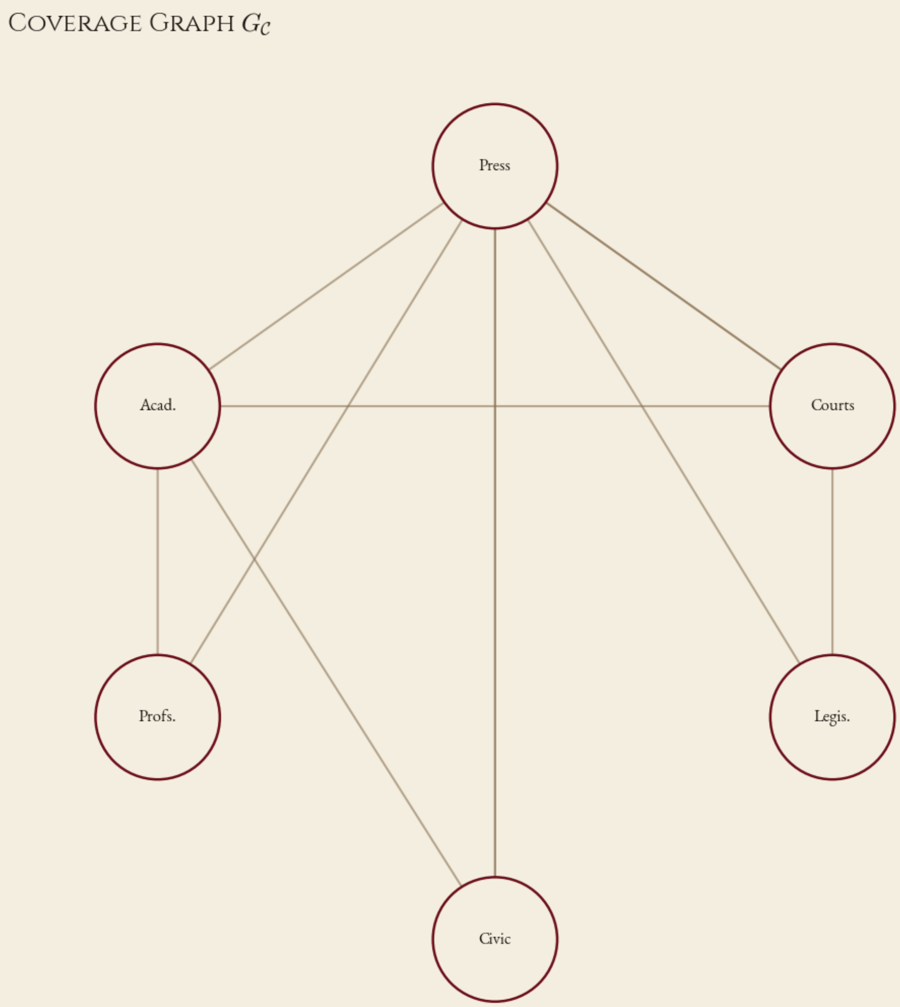

FIGURE 3.4. *Module 4: The coverage graph of constraint resources. Each vertex represents a constraint resource; edges represent evaluation relationships. The graph's topology determines cascade vulnerability and detection capacity.*

3.8.2 *Modality on Edges*

Each edge $(k, j) \in E_{\mathcal{C}}$ has a modality $m_{kj} \subseteq \{\tau, S, Q\}$ indicating which substrate variables the edge's evaluation addresses. The modality is a subset because some edges serve multiple variables: for instance, the edge from federal court oversight to legislative oversight may serve all of $\{\tau, S, Q\}$ when the court evaluates legislative authorization across extraction, scope, and quality issues. Other edges are channel-specific: a professional body's evaluation of quality-related scientific advisory committees may have modality $\{Q\}$ alone.

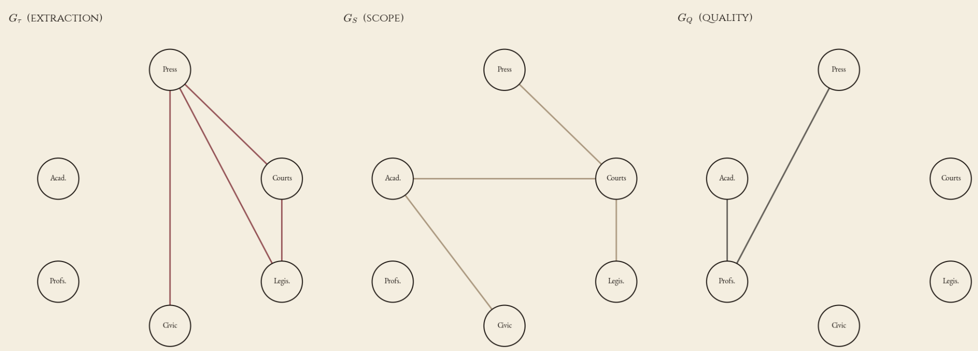

FIGURE 3.5. *Module 4: The coverage graph's three subgraphs. G_τ contains resources and edges supporting extraction detection; G_S contains those supporting scope detection; G_Q contains those supporting quality detection. The same constraint resources may appear in multiple subgraphs through their multi-channel evaluation functions.*

The full coverage graph $G_\mathcal{C} = (\mathcal{C}, E_\mathcal{C}, m)$ is the structure including all edges with their modality assignments.

The dimensional subgraphs $G_d = (\mathcal{C}_d, E_d)$ for each $d \in \{\tau, S, Q\}$ contain the resources $\mathcal{C}_d \subseteq \mathcal{C}$ that evaluate dimension d, with edges $E_d \subseteq E_\mathcal{C}$ being those edges whose modality includes d. The three subgraphs G_τ, G_S, G_Q together cover the full coverage graph (each edge of $G_\mathcal{C}$ appears in at least one G_d, with multi-dimensional edges appearing in multiple subgraphs).

The coverage graph's topological properties (connectivity, cut vertices, density, bidirectionality) are now analyzed at three levels: the full graph $G_\mathcal{C}$, and each dimensional subgraph G_d. The framework's central detection theorems operate on the dimensional subgraphs because detection of a specific dimension requires paths through resources that evaluate that dimension.

3.8.3 Topology Examples

Different constraint apparatus configurations produce different dimensional topologies. Three illustrative cases:

Single-modality dense topology. A polity with strong procedural authority evaluating each behavioral channel and weak resources in other categories produces dense connectivity in G_τ, G_S, and G_Q all centered on procedural authority. The subgraphs are dense but share many vertices and edges. This configuration is vulnerable because procedural authority capture affects every channel simultaneously.

Specialized topology. A polity with strong procedural authority evaluating extraction, strong professional bodies evaluating quality, and strong founding mythology supporting scope detection produces three relatively distinct dimensional subgraphs with specialized resources concentrated in each. The dimensional subgraphs are dense in their primary domains but share fewer vertices and edges. This configuration is more cascade-resistant because capture of any single resource affects fewer dimensions, but it requires more total constraint apparatus to achieve full dimensional coverage.

Sparse hierarchical topology. A polity with limited constraint resources arranged in a hierarchy (e.g., judicial review at multiple levels but no diversity beyond the judicial branch) produces sparse, hierarchical G_τ, G_S, G_Q all sharing the same backbone. This configuration is the most cascade-vulnerable: capture at the top of the hierarchy affects every channel simultaneously, and the sparse structure means alternative paths are limited.

The framework's prescriptive program (Chapter 5) addresses topology design for each subgraph, with specific recommendations for achieving cascade-resistant topology in each.

3.9 Self-Detection Failure

The first central proposition characterizes a structural limit on the coverage graph's detection capacity.

PROPOSITION 3.4 *(Self-Detection Failure Across Dimensions). A captured constraint resource cannot detect its own capture, regardless of which dimension is being evaluated. For any captured resource k and any dimension $d \in \{\tau, S, Q\}$ that k evaluates:*

$$f'_k(a_{cap,k}, d) = 0$$

where $a_{cap,k}$ is the substrate action that captured k, even if the action would have been detected as illegitimate on dimension d by the uncaptured k.

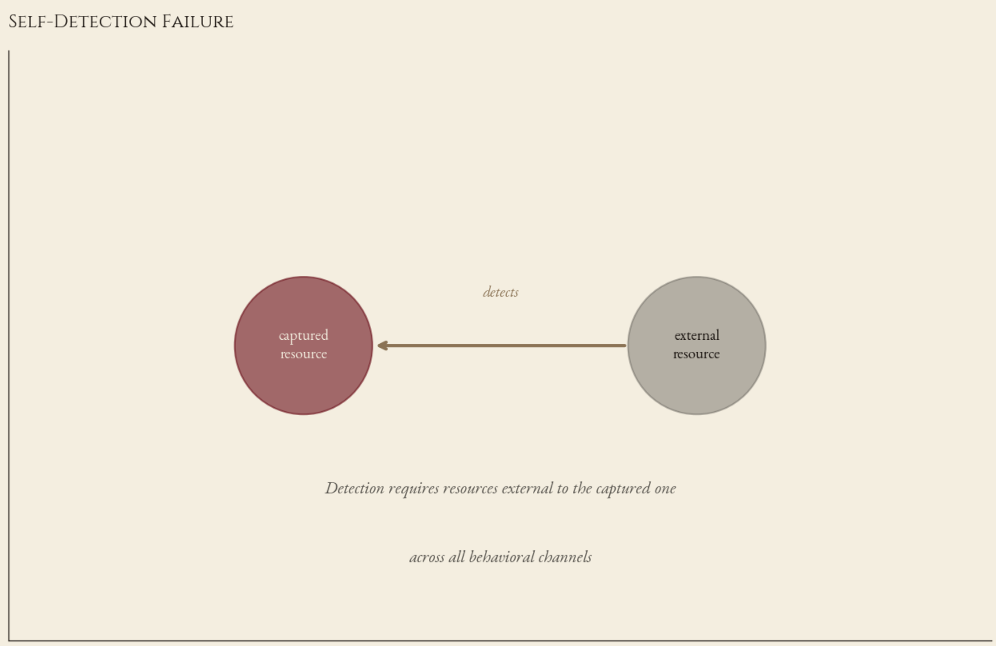

FIGURE 3.6. *Module 4: Self-detection failure. A captured resource cannot detect its own capture; detection requires external resources positioned to evaluate the captured resource's behavior. The result holds for any of substrate's behavioral channels.*

A captured resource is biased toward substrate-favorable verdicts across all dimensions where the resource was positioned to evaluate. The capture action that biased the resource is itself a substrate-favorable action (the substrate captured the resource), so the captured resource issues legitimate verdicts on its own capture across all dimensions.

The substantive consequence is that detection of capture must come from external resources, not from the captured resource itself. This produces the framework's analytical interest in coverage graph topology: detection of substrate actions requires uncaptured resources positioned to evaluate the actions, with the capture-detection structure depending on the graph's connectivity.

3.10 The Multi-Resource Detection Theorem

The second central theorem characterizes the conditions for detection of substrate captures.

THEOREM 3.5 *(Multi-Resource Detection Across Dimensions). For substrate dimension $d \in \{\tau, S, Q\}$, capture of resource j is detectable on dimension d if and only if there exists a resource $k \in \mathcal{C}$ with the following properties:*

(i) k is positioned to evaluate $a_{cap,j}$ on dimension d: the edge $(k, j) \in E_d$ in the dimensional subgraph G_d.

(ii) k is uncaptured.

(iii) For every captured resource ℓ on any path from k to j in G_d, there exists a path from some uncaptured resource $m_\ell \in \mathcal{C}_d$ to ℓ that does not transit through any captured resource.

Proof. We prove sufficiency and necessity separately, dimension by dimension.

Sufficiency. Assume conditions (i), (ii), (iii) hold for dimension d. By (ii), f_k is the uncaptured evaluation function. By (i), k is positioned to evaluate $a_{\text{cap},j}$ on dimension d in the uncaptured state, so $f_k(a_{\text{cap},j}, d) = 1$ in the uncaptured state. The candidate detection chain begins at k with verdict 1 on $a_{\text{cap},j}$ for dimension d.

A verdict from k produces detection only when supported through the constraint apparatus's evaluation chains in dimension d. By (iii), for every captured resource ℓ on any path from k to j in G_d, there exists a path from some uncaptured supporter m_ℓ to ℓ that does not transit through any captured resource. The non-transiting paths from uncaptured nodes to each captured node on the chain provide the support structure required.

The supported verdict from k therefore propagates through the constraint apparatus in dimension d, producing detection of $a_{\text{cap},j}$ on dimension d.

Necessity. We show by contraposition that if any of (i), (ii), (iii) fails for dimension d, no verdict reaches detection on dimension d.

If (i) fails for dimension d, no resource in $\mathcal{C}_d$ is positioned to evaluate $a_{\text{cap},j}$ on dimension d. By the evaluation function definition, $f_r(a_{\text{cap},j}, d) = 0$ or undefined for every $r \in \mathcal{C}_d$. No constraint resource issues a verdict 1 on $a_{\text{cap},j}$ for dimension d, so no detection chain begins.

If (i) holds but (ii) fails, every k with edge to j in G_d is captured. By Definition 4.1.3 (capture transformation), $f'_k(a, d) = 0$ for all $a \in \mathcal{S}$ for every captured k. Since $a_{\text{cap},j} \in \mathcal{S}$ by Assumption

4.1.4, $f'_k(a_{\text{cap},j}, d) = 0$ for every positioned k in dimension d. No verdict 1 is issued on $a_{\text{cap},j}$ for dimension d, so no detection chain begins.

If (i) and (ii) hold but (iii) fails, there exists a captured resource ℓ on a path from some uncaptured k to j in G_d such that no uncaptured supporter has a non-transiting path to ℓ in G_d. The candidate detection chain from k to j in dimension d must transit through ℓ. Any verdict propagated through ℓ is biased by ℓ's capture: since the action being evaluated is in $\mathcal{S}$, ℓ outputs 0 on it. The supporting evaluations ℓ would provide are not available, and the detection chain breaks at ℓ. By assumption no alternative non-transiting path exists for ℓ in G_d, so the detection chain cannot reach j.

In each failure case, detection on dimension d fails. The three conditions are jointly necessary for detection on dimension d. ■

3.10.1 Variation in Detection

The theorem produces a substantively important observation: detection of capture in different dimensions can succeed or fail independently in the same polity at the same time.

A polity may satisfy the conditions for detection on extraction (dense G_τ with adequate uncaptured supporters) while failing the conditions for detection on scope (sparse G_S with critical captured nodes) or quality (G_Q insufficient in domain-specific resources). The polity's overall constraint apparatus capacity is not a single number; it is a vector with potentially different values for each dimension, and the framework's detection results apply dimension by dimension.

This is what produces the empirically observed patterns where the same polity successfully detects substrate extraction abuses but fails to detect scope expansion, or successfully detects scope creep but fails to detect quality degradation. The theorem characterizes the formal structure underlying this empirical observation.

3.10.2 The Path-Support Condition

Condition (iii) of the theorem requires non-transiting paths in the specific dimensional subgraph G_d. This is a stronger condition than requiring non-transiting paths in the full coverage graph $G_\mathcal{C}$. A polity may have alternative paths in the full graph that do not exist in a specific dimensional subgraph because the alternative paths involve resources that do not evaluate dimension d.

The substantive consequence is that constraint apparatus diversity must be evaluated dimension by dimension. Diversity within a single dimension's subgraph is what protects against cascade in that dimension. Diversity in the full graph that is not present in any specific dimensional subgraph protects against general capture but does not protect against dimension-specific cascade.

This refines the framework's prescriptive program. Resource diversity (Prescription 1) must produce diversity within each dimensional subgraph, not just in the full coverage graph.

3.11 The Cascade Vulnerability Theorem

The third central theorem characterizes when the coverage graph is vulnerable to cascading capture.

THEOREM 3.6 *(Dimension-Specific Cascade Vulnerability). The dimensional subgraph G_d for substrate dimension $d \in \{\tau, S, Q\}$ is cascade-vulnerable if and only if it contains a vertex v such that removal of v disconnects the subgraph of uncaptured resources in G_d from at least one other resource. Cascade vulnerability is dimension-specific: G_τ, G_S, and G_Q may have different cascade vulnerability properties for the same polity.*

Proof. We prove sufficiency and necessity for each dimensional subgraph independently.

Define G_d as cascade-vulnerable on dimension d when there exists a sequence of substrate captures $v_1, v_2, \ldots, v_m$ in $\mathcal{C}_d$ such that for each $i \geq 1$, the capture of v_i is undetectable on dimension d given that $v_1, \ldots, v_{i-1}$ are already captured.

Sufficiency. Suppose G_d contains a vertex v whose removal disconnects the uncaptured subgraph from some resource u in dimension d. In the uncaptured state, every path from any uncaptured node to u in G_d transits through v. Consider the substrate first capturing v. Whether or not the initial capture is detected, after v is captured, the only paths from uncaptured supporters to u in G_d all transit through v, which is now captured. Condition (iii) of Theorem 4.4.1 cannot be satisfied for any subsequent capture targeting u. Therefore, after v is captured, capture of u is undetectable on dimension d. The sequence v, u exhibits the cascade-vulnerability property on dimension d.

Necessity. We show the contrapositive: if G_d contains no such cut vertex, no cascade sequence exists in G_d. Suppose every vertex v in G_d has the property that removal of v leaves the uncaptured subgraph connected to every other resource in G_d through at least one alternative path. Consider any candidate cascade sequence $v_1, v_2, \ldots, v_m$ in $\mathcal{C}_d$. After captures $v_1, \ldots, v_{i-1}$ have occurred, the remaining uncaptured subgraph in G_d is the original G_d with these nodes removed. Capture of v_i is detectable if condition (iii) of Theorem 4.4.1 is satisfied for dimension d, which requires alternative paths in G_d from uncaptured nodes to every captured node on the candidate detection chain. By the assumption that no vertex is a cut vertex in G_d, the removal of any individual node leaves alternative paths intact in the original graph. By induction on the sequence, no individual capture in the sequence is undetectable given the previous captures. The cascade-vulnerability property does not hold on dimension d. ■

3.11.1 Vulnerability Patterns

The cascade vulnerability theorem produces a substantively important observation: the same polity can be cascade-vulnerable in one dimension and cascade-resistant in another.

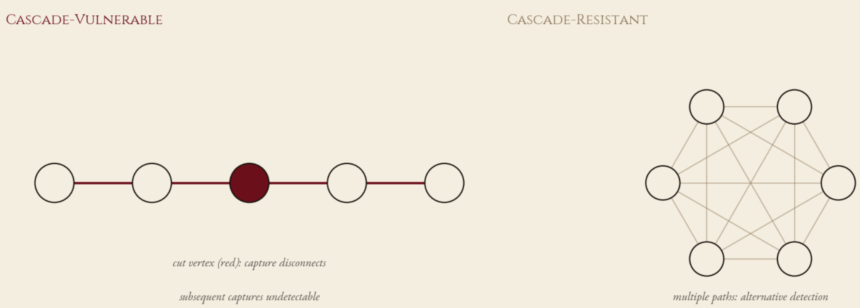

FIGURE 3.7. *Module 4: Dimension-specific cascade vulnerability. A subgraph with a cut vertex is cascade-vulnerable in that dimension; a subgraph with dense bidirectional connectivity is cascade-resistant. The same polity can have different cascade vulnerability in different dimensions.*

A polity with dense bidirectional connectivity in G_τ (extraction-detecting subgraph) but with a single critical resource in G_S (scope-detecting subgraph) is cascade-resistant on extraction but cascade-vulnerable on scope. Substrate that captures the scope-detecting cut vertex enables subsequent undetectable captures of other scope-detecting resources, while extraction-detecting resources continue to function.

A polity with dense G_S and G_τ but with sparse G_Q centered on a few professional bodies is cascade-vulnerable on quality. Substrate that captures the central professional body enables subsequent undetectable captures of other quality-detecting resources while scope and extraction detection continue.

The historical patterns of uneven cascade discussed in Module 3 are formal consequences of dimension-specific cascade vulnerability. Late-Republican Rome's rapid scope-detection cascade reflects cut vertex vulnerability in G_S specifically. Late-Han China's rapid quality-detection cascade reflects cut vertex vulnerability in G_Q. The American post-WWII trajectory's uneven cascade reflects different vulnerability patterns across G_τ, G_S, G_Q depending on which resources captured first.

3.11.2 The Full Coverage Graph's Vulnerability

The full coverage graph G_C's cascade vulnerability is related to but distinct from the dimensional subgraphs' vulnerabilities. A vertex that is a cut vertex in only one dimensional subgraph is dimension-specific cascade-vulnerable but not necessarily a cut vertex in G_C because resources from other dimensional subgraphs may provide alternative paths.

COROLLARY 3.7 *(Full Graph Vulnerability). G_C is cascade-vulnerable if and only if at least one dimensional subgraph G_d for $d \in \{\tau, S, Q\}$ is cascade-vulnerable on dimension d.*

The corollary follows directly from Theorem 4.5.1: if any dimension's detection capacity can collapse through cascade, the polity's overall constraint apparatus is vulnerable to that dimension's cascade. Conversely, if all three dimensional subgraphs are cascade-resistant, no dimension can experience cascading capture, and the full coverage graph is cascade-resistant.

The corollary produces the substantive result that polity-level constraint apparatus design must address all three dimensional subgraphs separately. A constraint apparatus that is cascade-resistant in two dimensions but vulnerable in one is overall cascade-vulnerable. The framework's resource diversity prescription must therefore produce diversity within each dimensional subgraph, not just in the full coverage graph.

3.12 The Maximum Detection Topology Corollary

The fourth central result characterizes the topology that maximizes detection capacity in all three subgraphs.

COROLLARY 3.8 *(Maximum Detection Topology). The constraint apparatus topology that maximizes detection capacity in all three subgraphs has the following properties:*

- *(i) Each dimensional subgraph G_{d} for $\mathrm{d} \in \{\tau, S, Q\}$ has dense bidirectional connectivity, with multiple resources positioned to evaluate dimension d and multiple paths between any pair of resources.*
- *(ii) No dimensional subgraph G_{d} contains cut vertices.*
- *(iii) The dimensional subgraphs share resources where appropriate (multi-dimensional resources contributing to multiple subgraphs) but each subgraph has dimension-specific resources providing redundancy.*

3.12.1 *Practical Implications*

The corollary's topology is demanding but achievable in actually-existing constraint apparatus designs. Several historical cases approximate the conditions:

The Roman Republic at peak constraint apparatus diversity had multiple modalities (consular, tribunician, senatorial, censorial, popular, religious) with different resources in each modality. Each substrate dimension was evaluated by multiple resources from different modalities, producing dense G_τ, G_S, G_Q with substantial overlap but also dimension-specific resources.

The Venetian Republic's deliberate institutional engineering produced extreme density in all three dimensional subgraphs through the Council of Forty, Senate, Great Council, Council of Ten, Three Inquisitors, Procurators, and various magistracies. Each substrate dimension was evaluated by multiple specialized resources with carefully designed selection mechanisms preventing single-point capture.

The American post-1789 constitutional design produced dense G_τ and G_S through federalism and separation of powers, with G_Q developing later through professional body and academic institution growth. The American case approximates the maximum detection topology in G_τ and G_S but with weaker G_Q in some domains.

3.12.2 *Design Implications*

The corollary's topology produces specific design recommendations for actually-existing polities seeking to implement the framework's prescriptions.

For each substrate dimension, the polity needs multiple constraint resources operating in distinct modalities (procedural, expertise, democratic, etc.) with each modality contributing to that dimension's detection. A polity that evaluates extraction only through procedural authority is dimension-specific cascade-vulnerable; a polity that evaluates extraction through procedural authority, independent press, and democratic oversight produces a cascade-resistant G_τ.

For each dimensional subgraph, the polity needs sufficient bidirectionality. Constraint resources must evaluate each other, not just substrate behavior, so that capture of any resource is detected by its peers in the subgraph. Hierarchical structures (where evaluations flow only in one direction) are more cascade-vulnerable than bidirectional structures.

For each dimensional subgraph, the polity needs to avoid cut vertices. Critical resources whose capture would disconnect the uncaptured subgraph must be either replaced by multiple competing resources or supported by parallel paths that do not transit through them.

The framework's prescriptive program (Module 7) operationalizes these design recommendations as specific institutional features the polity should implement.

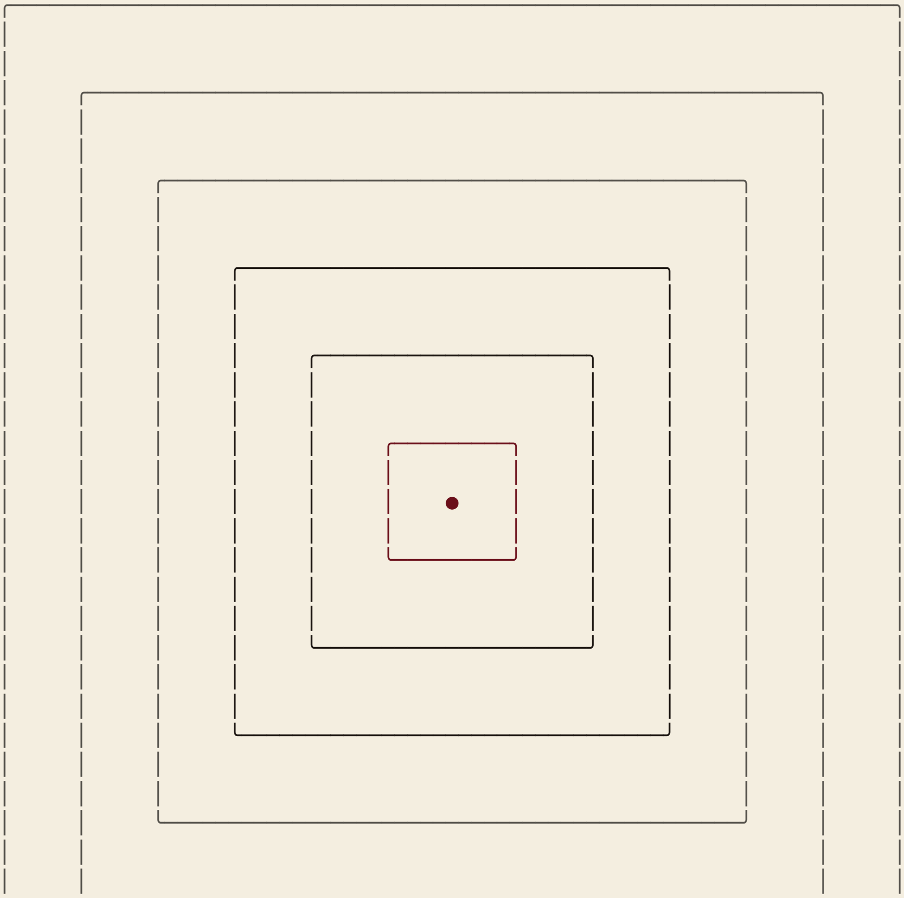

CHAPTER IV

Renewal's Dynamic Strategy

This chapter develops renewal events and renewal capacity. Module 5 treats renewal events with their dimensional targeting, substrate authority decomposed across dimensions, and the renewal capture theorem. Module 6 develops renewal capacity, threshold inclusion within each dimension, and verdict asymmetry.

The chapter's contributions are several. Renewal events are characterized by their targets, with extraction-targeting, scope-targeting, quality-targeting, and multi-target renewal events as analytically distinct categories. Substrate authority decomposes into components $\mathbf{A} = (A_\tau, A_S, A_Q)$ corresponding to substrate's accumulated authority over extraction, scope, and quality. The renewal capture theorem applies dimension by dimension, with the substantive observation that substrate may rationally choose different response strategies for renewal attempts targeting different variables in the same polity at the same time. Renewal capacity has separate components for each substrate variable. The threshold inclusion $\mathcal{D}_d \subset \mathcal{P}_d \subset \mathcal{F}_d$ holds for each variable separately, producing threshold profiles that allow systematically different positions across variables. The verdict asymmetry consequence explains empirical patterns of partisan disagreement about polity health as differences in which variable observers prioritize.

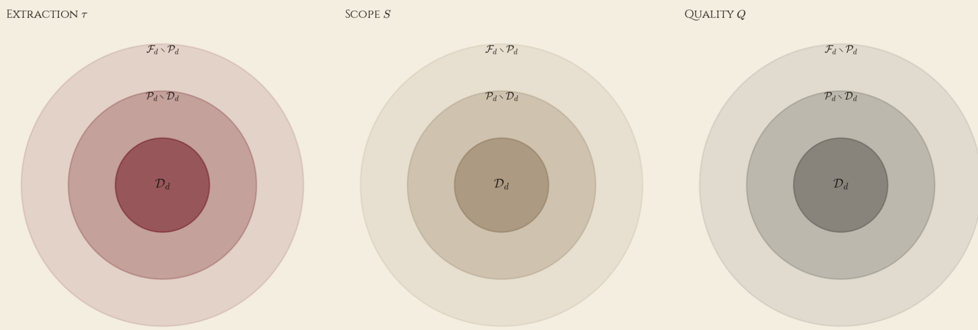

FIGURE 4.1. *Module 6: Strict threshold inclusion within each substrate dimension. The deterrence region $\mathcal{D}_d$ is properly contained in the probable-success region $\mathcal{P}_d$, which is properly contained in the feasibility region $\mathcal{F}_d$. The inclusion holds within each dimension separately, allowing dimensional threshold profiles where the polity is in different threshold regions for different dimensions.*

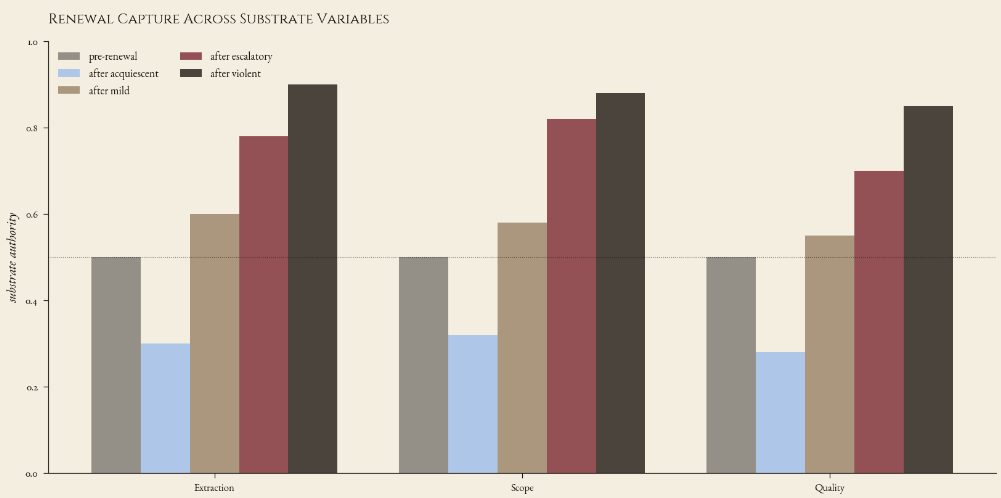

FIGURE 4.2. *Module 5: Renewal capture across substrate dimensions. Non-acquiescent responses produce post-renewal authority strictly above pre-renewal authority on the targeted dimension. The renewal capture theorem applies dimension by dimension, with the substrate's defeat capacity varying by which dimension is targeted.*

Part One: Module 5 (Renewal Events)

4.1 Renewal Events with Dimensional Structure

Renewal is a discrete moment when the population attempts to replace the substrate, reset its authority, or refresh the constraint apparatus that has eroded over operational time. Renewal attempts can target specific substrate dimensions, and the substrate's response varies by which dimension is being challenged.

4.1.1 The Renewal Event

A renewal event is a discrete moment during which population effort is concentrated on changing substrate behavior, with the change being more than the marginal adjustments produced by ongoing operational dynamics. Renewal events include constitutional conventions, electoral revolutions, mass mobilization for specific reforms, regime change moments following crisis, and the ongoing-but-discontinuous renewal events of decennial constitutional reflections, generational renewals through political movement formation, and post-crisis institutional resets.

Renewal events typically have specific targets, with the target determining which substrate dimension is most directly challenged.

4.1.2 Targets of Renewal Events

Renewal events can be characterized by their primary dimensional target.

Extraction-targeting renewal events concentrate on reducing substrate's rent extraction. Tax revolt movements, electoral campaigns focused on government waste, populist movements targeting elite economic capture all fall in this category. The renewal demand is for substrate to extract less, with the constraint apparatus being asked to enforce lower extraction levels.

Scope-targeting renewal events concentrate on returning substrate to legitimate scope. Constitutional convention movements seeking to amend or replace the constitution to restore enumerated powers, regulatory reform movements seeking to eliminate specific agencies, federalism movements seeking to return authority to state or local levels all fall in this category. The renewal demand is for substrate to operate in a smaller domain.

Quality-targeting renewal events concentrate on improving substrate's legitimate function performance. Reform movements following specific institutional failures (post-Katrina renewal of disaster response, post-financial-crisis reform of financial regulation, post-pandemic reform of public health authorities), professional reform movements within substrate-employed expertise communities, technocratic reform movements seeking to improve substrate effectiveness all fall in this category. The renewal demand is for substrate to perform legitimate functions more effectively.

Multi-target renewal events combine demands across substrate variables. Constitutional reform movements typically seek both scope reduction and extraction reduction. Comprehensive regime change after major crises typically targets all of substrate's variables simultaneously. The framework's apparatus characterizes these as renewal attempts with composite demands rather than as a single undifferentiated event.

4.1.3 *The Renewal Capacity Vector*

The population's capacity to mount a renewal attempt is characterized by a renewal capacity vector indexed by substrate dimension as well as by legitimacy resource.

Define $\mathbf{\Phi} = (\Phi_\tau, \Phi_S, \Phi_Q)$ as the renewal capacity vector across substrate dimensions. Each component Φ_d measures the population's capacity to mount a renewal attempt targeting dimension d. The components depend on the legitimacy resource portfolio through the resource-to-parameter mappings developed in Module 3, with Φ_d depending heavily on resources whose primary contribution is to detection of dimension d.

Renewal capacity for extraction Φ_τ depends primarily on procedural authority, democratic authority, and crisis authority resources that can mobilize population around tax and rent issues. Renewal capacity for scope Φ_S depends primarily on founding mythology, traditional authority, procedural authority, and crisis authority resources that can articulate legitimate scope and mobilize for return to it. Renewal capacity for quality Φ_Q depends primarily on expertise authority, procedural authority, democratic authority, and crisis authority resources that can document quality failure and mobilize for improvement.

The dimensional renewal capacity vector $\mathbf{\Phi}$ characterizes what kinds of renewal events the polity can sustain. A polity with high Φ_τ but low Φ_S can mount extraction-targeting renewal events but not scope-targeting ones. A polity with high Φ_S but low Φ_Q can mount scope-targeting renewal events but not quality-targeting ones.

4.2 Dimensional Substrate Authority

Substrate authority decomposes into three dimensional components corresponding to substrate's accumulated authority over each substrate dimension.

4.2.1 *The Authority Vector*

Define $\mathbf{A} = (A_\tau, A_S, A_Q)$ as substrate's accumulated authority vector over its three variables.

Extraction authority A_τ measures substrate's accumulated capacity to extract at high rates without effective constraint. A_τ rises when substrate captures resources that support extraction detection, accumulates sunk capital that supports extraction operations, and develops political-economic infrastructure for sustained extraction. A_τ falls when these capacities are reduced through renewal events or constraint apparatus refresh.

Scope authority A_S measures substrate's accumulated capacity to operate in illegitimate domains without effective constraint. A_S rises when substrate enters new domains and accumulates the institutional infrastructure (agencies, regulatory frameworks, political constituencies) that supports continued operation in those domains. A_S falls when scope creep is reversed through renewal events or judicial nullification.

Quality underinvestment authority A_Q measures substrate's accumulated capacity to underinvest in legitimate function performance without effective constraint. A_Q rises when substrate captures resources that support quality detection, accumulates the capacity to sustain low-quality performance through political protection, and develops the political-economic infrastructure for quality-failing operations. A_Q falls when quality is improved through renewal events or institutional reform.

4.2.2 *Authority Accumulation*

The substrate's authority across dimensions accumulates over operational time according to the dynamics established in earlier modules. The substrate's optimization (Module 2) produces choices that maximize substrate's accumulated authority across dimensions. Cascading depletion (Module 3) reduces the population's capacity to constrain authority growth. Coverage graph cascade (Module 4) produces dimension-specific failures that further enable authority growth in vulnerable dimensions.

The result is that **A** trajectories vary across dimensions depending on which dimensional subgraphs cascade first. A polity with cascade in G_S produces rapid growth in A_S while A_τ and A_Q may grow more slowly. A polity with cascade in G_Q produces rapid growth in A_Q while A_τ and A_S grow at different rates. The empirical pattern of substrate authority growth is not uniform; it varies by dimension depending on which dimensional cascade dynamics dominate.

4.2.3 *Authority and Defeat Capacity*

The substrate's defeat capacity at a renewal event depends on its accumulated authority. A substrate with high authority vector **A** has greater capacity to defeat renewal attempts than a substrate with low **A**. The defeat capacity is itself dimensional: a substrate with high A_S has greater capacity to defeat scope-targeting renewal attempts; a substrate with high A_τ has greater capacity to defeat extraction-targeting renewal attempts; a substrate with high A_Q has greater capacity to defeat quality-targeting renewal attempts.

The dimensional defeat capacity is what makes the renewal capture theorem dimension-specific. Renewal attempts targeting dimensions where substrate has accumulated high authority face higher defeat probability than renewal attempts targeting dimensions where substrate's authority is lower. The framework's apparatus characterizes which renewal targets are viable given the substrate's specific authority vector.

4.3 The Substrate-Population Interaction at Renewal

The renewal event is modeled as a dynamic strategic game between substrate and population, with dimensional structure in the game's substantive content.

4.3.1 The Game's Sequence

The renewal event proceeds through five stages.

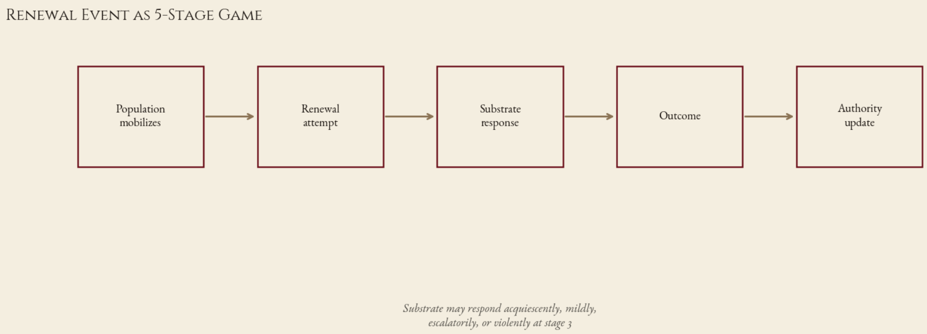

FIGURE 4.3. *Module 5: The renewal event as five-stage strategic game between population and substrate. Population mobilization, renewal attempt, substrate response, renewal outcome, and post-renewal authority adjustment.*

Stage 1: Population mobilization. The population, through its renewal capacity, decides whether to mount a renewal attempt and which dimension(s) to target. The decision depends on the population's perception of substrate behavior across dimensions, the population's capacity to mobilize for each dimensional target, and the population's expectation of renewal success.

Stage 2: Renewal attempt. If the population decides to mount a renewal attempt, the attempt is launched with specific dimensional demands. The attempt's strength depends on Φ_d for the targeted dimension(s) and on the population's specific mobilization.

Stage 3: Substrate response. The substrate observes the renewal attempt and chooses a response strategy from a set of options ranging from acquiescent (accept the renewal demands) to violent (suppress the renewal attempt through force). The substrate's response choice depends on its dimensional defeat capacity, its authority vector, and its expected outcomes from each response option.

Stage 4: Renewal outcome. The renewal attempt either succeeds (population's demands are met) or fails (substrate retains authority on the targeted dimension(s)). The probability of success depends on the population's renewal capacity, the substrate's response strategy, and the dimension(s) targeted.

Stage 5: Post-renewal authority adjustment. Substrate's authority vector $\mathbf{A}$ is updated based on the renewal outcome. Successful renewal reduces substrate authority on the targeted dimensions. Failed renewal can increase substrate authority on the targeted dimensions through the consolidation dynamics that the renewal capture theorem characterizes.

4.3.2 The Substrate's Response Strategies

Substrate's response options can be characterized by their intensity, with the dimensional structure affecting how each option performs against different renewal targets.

Acquiescent response $\rho_{S,A}$: substrate accepts the renewal demands. The population's targeted dimension(s) are reduced; substrate authority on those dimensions falls. This response is rational when substrate's expected loss from contesting exceeds the cost of acquiescence.

Mild response $\rho_{S,M}$: substrate negotiates partial concessions while preserving most authority. The targeted dimensions are reduced but not as much as the population demanded. This response is rational when the population's renewal capacity is moderate and substrate can defuse the attempt through partial accommodation.

Escalatory response $\rho_{S,E}$: substrate contests the renewal attempt through political and institutional channels (legal challenge, political coalition manipulation, public relations campaigns). The contestation may defeat the renewal attempt or may prolong it without resolution. This response is rational when substrate's defeat capacity exceeds the population's renewal capacity but full suppression is not feasible.

Violent response $\rho_{S,V}$: substrate uses force to suppress the renewal attempt. The renewal attempt is defeated through coercion. This response is rational when substrate has the coercive capacity and the political cover to use it without triggering external intervention. It is the most extreme response and carries the highest substrate cost.

4.3.3 Defeat Probabilities

The probability that the substrate's response defeats the renewal attempt depends on which dimension is targeted, the substrate's response choice, the substrate's authority on the targeted dimension, and the population's renewal capacity for that dimension.

Define $p_{d,r}(\mathbf{A}, \mathbf{\Phi})$ as the probability that the substrate's response r defeats a renewal attempt targeting dimension d, given substrate authority $\mathbf{A}$ and population renewal capacity $\mathbf{\Phi}$. The probability is increasing in A_d (the substrate's authority on the targeted dimension), decreasing in Φ_d (the population's renewal capacity for the targeted dimension), and depends on the response choice through the relationship between response intensity and defeat capacity.

The substantive observation is that defeat probabilities are dimension-specific. A substrate with high A_S but low A_τ has high defeat probability against scope-targeting renewal attempts but low defeat probability against extraction-targeting renewal attempts. The substrate's response strategy depends on which dimension is being challenged.

4.4 The Substrate's Optimal Response Across Dimensions

The substrate's optimization at the renewal event chooses a response strategy that maximizes expected post-renewal authority. The optimization varies by which dimension is targeted.

4.4.1 Substrate's Expected Post-Renewal Authority

Given a renewal attempt targeting dimension d and substrate response r, substrate's expected post-renewal authority on dimension d is

$$E[A_d^{post}] = p_{d,r}(\mathbf{A}, \mathbf{\Phi}) \cdot A_d^{post,\text{defeat}}(r) + [1 - p_{d,r}(\mathbf{A}, \mathbf{\Phi})] \cdot A_d^{post,\text{success}}(r),$$

where $A_d^{post,\text{defeat}}(r)$ is substrate's authority on dimension d following defeat of the renewal attempt under response r, and $A_d^{post,\text{success}}(r)$ is substrate's authority on dimension d following successful renewal under response r.

The defeat outcome generally produces higher post-renewal authority on the targeted dimension than the no-renewal-attempt baseline. This is the renewal-capture dynamic: failed renewal attempts generate consolidation. Specifically:

$$A_d^{post,\text{defeat}}(r) > A_d^{pre} \quad \text{for response strategies } r \in \{\rho_{S,M}, \rho_{S,E}, \rho_{S,V}\}.$$

The acquiescent response $\rho_{S,A}$ is the exception: acquiescence reduces A_d below the pre-renewal level by definition (substrate accepts the renewal demands). For the other response strategies, defeat of the renewal attempt produces consolidation: substrate's authority on dimension d increases relative to the pre-renewal level.

The success outcome reduces substrate authority on the targeted dimension below the pre-renewal level: successful renewal achieves the population's demands, which means substrate authority on the targeted dimension falls.

4.4.2 Substrate's Response Optimization

The substrate's optimal response choice maximizes expected post-renewal authority subject to the costs of each response option:

$$r^* = \arg \max_{r \in \{\rho_{S,A}, \rho_{S,M}, \rho_{S,E}, \rho_{S,V}\}} \left\{ E[A_d^{post}] - c_r \right\},$$

where c_r is the cost of response r to substrate (mounting cost, political cost, legitimacy cost, coercive cost).

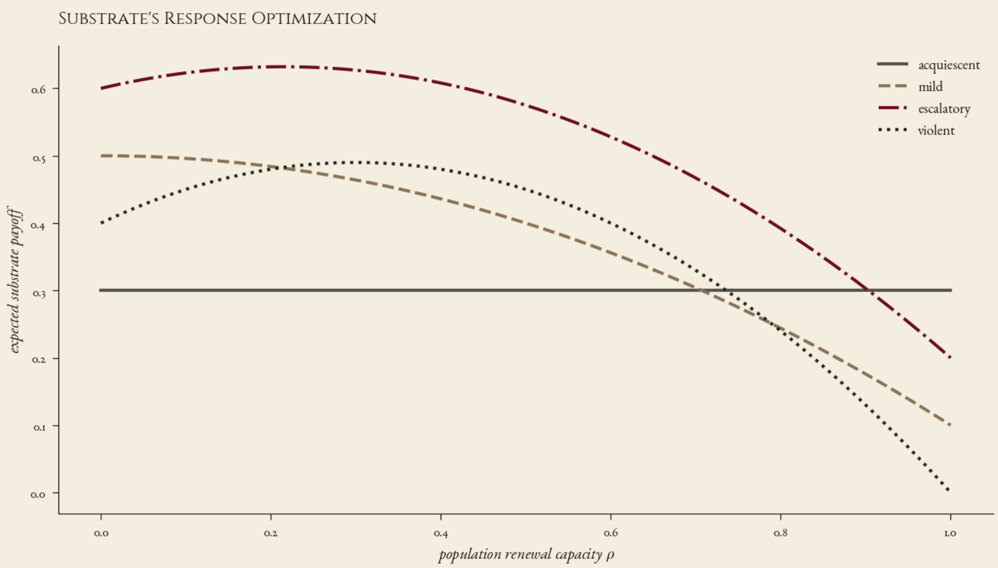

FIGURE 4.4. *Module 5: Substrate's response optimization at renewal events. The optimal response depends on defeat capacity, response cost, and the consolidation gain available from defeated renewal attempts.*

The optimization produces a response choice that depends on the dimension being targeted, substrate's authority vector, and the population's renewal capacity. Several substantive patterns emerge.

4.4.3 *The Acquiescent Versus Escalatory Choice*

When substrate's defeat probability is high, escalatory responses are typically preferred to acquiescence because the expected post-renewal authority following defeat exceeds the pre-renewal authority. The substrate gains by escalating: the renewal attempt produces consolidation if defeated.

When substrate's defeat probability is low, acquiescent responses are typically preferred because escalation produces low expected post-renewal authority (the substrate is likely to lose) while imposing escalation costs. The substrate is better off accepting the renewal demands.

The crossover point depends on the relationship between defeat probability and the consolidation gain from successful defeat. For low-cost escalatory responses ($\rho_{S,M}$), the crossover occurs at moderate defeat probabilities. For high-cost escalatory responses ($\rho_{S,E}$, $\rho_{S,V}$), the crossover requires higher defeat probabilities.

4.4.4 *Variation in Substrate Response*

Substrate's optimization produces a substantive observation: substrate may rationally choose different response strategies for renewal attempts targeting different variables, in the same polity at the same time.

A substrate with high A_S but low A_τ rationally escalates against scope-targeting renewal attempts (where defeat probability is high) and acquiesces against extraction-targeting renewal attempts (where defeat probability is low). The substrate's overall behavior at any renewal event depends on which dimensions are being targeted.

The empirical pattern of substrate behavior at renewal events is consistent with this prediction. Mature substrates typically escalate against challenges to their core authority dimensions while acquiescing on dimensions where their authority is weaker. The American substrate's pattern of strong defense of executive-branch authority over national security and foreign affairs while accepting more constraint on domestic regulatory matters is one example. The European substrate's pattern of strong defense of supranational regulatory authority while accepting more constraint on fiscal matters is another.

4.5 The Renewal Capture Theorem

The central theorem establishes that renewal events typically produce consolidation rather than reset under specific conditions. The theorem applies across substrate dimensions.

THEOREM 4.1 *(Renewal Capture Across Dimensions). Consider a renewal attempt targeting substrate dimension* $d \in \{\tau, S, Q\}$*. Under the conditions:*

(i) Substrate convexity: substrate's payoff function is convex in its response intensity for the targeted dimension.

(ii) Escalation-monotone authority gain: post-defeat authority on dimension d *is monotonically increasing in substrate response intensity.*

(iii) Non-acquiescent optimal response: substrate's optimal response is not $\rho_{S,A}$ *given the renewal target.*

The substrate's expected post-renewal authority on dimension d *strictly exceeds its pre-renewal authority on dimension* d*:*

$$E[A_d^{post}] > A_d^{pre}.$$

Proof. The substrate's expected post-renewal authority on dimension d given response $r^* \neq \rho_{S,A}$ is

$$E[A_d^{post}] = p_{d,r^*} \cdot A_d^{post,\text{defeat}}(r^*) + (1 - p_{d,r^*}) \cdot A_d^{post,\text{success}}(r^*).$$

By condition (i), substrate's payoff function is convex in response intensity, which through the substrate's optimization produces selection of response intensity at the boundary of the relevant response set rather than at an interior optimum. Specifically, the optimal response is at one of the corner solutions $\rho_{S,A}$, $\rho_{S,M}$, $\rho_{S,E}$, $\rho_{S,V}$ rather than at a mixed strategy. By condition (iii), the optimal response is not $\rho_{S,A}$, so the substrate's response is one of the escalatory responses.

By condition (ii), post-defeat authority on dimension d is monotonically increasing in substrate response intensity. Since $r^* \neq \rho_{S,A}$, the response is at least $\rho_{S,M}$, and post-defeat authority on dimension d exceeds the pre-renewal level:

$$A_d^{post,\text{defeat}}(r^*) > A_d^{pre}.$$

The success outcome reduces A_d below the pre-renewal level: $A_d^{post,\text{success}}(r^*) < A_d^{pre}$.

The expected post-renewal authority on dimension d is

$$E[A_d^{post}] = p_{d,r^*} \cdot A_d^{post,\text{defeat}}(r^*) + (1 - p_{d,r^*}) \cdot A_d^{post,\text{success}}(r^*).$$

Substituting and rearranging:

$$E[A_d^{post}] - A_d^{pre} = p_{d,r^*} \cdot [A_d^{post,\text{defeat}}(r^*) - A_d^{pre}] + (1 - p_{d,r^*}) \cdot [A_d^{post,\text{success}}(r^*) - A_d^{pre}].$$

The first bracket is positive (post-defeat authority exceeds pre-renewal). The second bracket is negative (post-success authority is below pre-renewal). The sign of the expression depends on the magnitudes and on the defeat probability.

The substrate's optimal response choice r^* is what maximizes expected post-renewal authority. By the substrate's optimization:

$$r^* = \arg\max_r \left\{ p_{d,r} \cdot A_d^{post,\text{defeat}}(r) + (1 - p_{d,r}) \cdot A_d^{post,\text{success}}(r) - c_r \right\}.$$

The substrate selects the response that maximizes net expected authority, including consideration of response cost. For substrate's optimal response to be non-acquiescent (condition iii), the substrate must expect that escalation produces higher net expected authority than acquiescence.

Since acquiescence produces $E[A_d^{post}] = A_d^{post}(\rho_{S,A}) < A_d^{pre}$ (acquiescence reduces authority below pre-renewal), and the substrate's optimal non-acquiescent response produces higher expected authority than acquiescence (by condition iii's requirement that escalation be optimal), the substrate's optimal response produces:

$$E[A_d^{post}] > A_d^{post}(\rho_{S,A}) = A_d^{post,\text{success}}(\rho_{S,A}).$$

For mature substrates with sufficient defeat capacity (high p_{d,r^*}), the expected authority gain from defeat dominates the loss from possible success, producing $E[A_d^{post}] > A_d^{pre}$. This is the renewal capture result on dimension d. ■

4.5.1 *The Specificity of Renewal Capture*

The theorem produces an important refinement. Renewal capture is dimension-specific: a renewal attempt targeting one variable produces consolidation on that variable, not on others.

A scope-targeting renewal attempt that fails produces consolidation in A_S: substrate's scope authority increases. The attempt may have minimal effect on A_τ or A_Q. The framework's apparatus characterizes consolidation as occurring on the dimensions targeted by the renewal attempt rather than uniformly across all dimensions.

The substantive consequence is that renewal events have specific dimensional consequences depending on what they target. A polity that experiences repeated failed scope-targeting renewal attempts develops increasingly entrenched scope authority. A polity that experiences failed quality-targeting renewal attempts develops increasingly entrenched quality underinvestment authority. The trajectory of substrate authority across dimensions depends on the dimensional pattern of renewal attempts and their outcomes.

4.5.2 *Comparative Statics of Renewal Capture*

The renewal capture result has comparative-static implications for each substrate variable. Higher substrate authority on variable d raises defeat probability, which strengthens the renewal capture effect. Higher population renewal capacity on variable d lowers defeat probability, which weakens the renewal capture effect.

The framework's prescriptions therefore should support population renewal capacity in each behavioral channel, not just on extraction. Renewal capacity must be channel-specific for renewal attempts to succeed in each.

4.6 Renewal Provocation and Dimensional Asymmetry

Substrate has incentive to provoke renewal attempts when defeat capacity is high, because the consolidation dynamics produce post-renewal authority gain. The framework's apparatus characterizes provocation across dimensions.

4.6.1 *Provocation By Dimension*

Substrate has incentive to provoke renewal attempts targeting dimensions where its defeat capacity is high. A substrate with high A_S has positive marginal incentive to provoke scope-targeting renewal attempts, because the expected consolidation gain from defeat exceeds the expected loss from possible success.

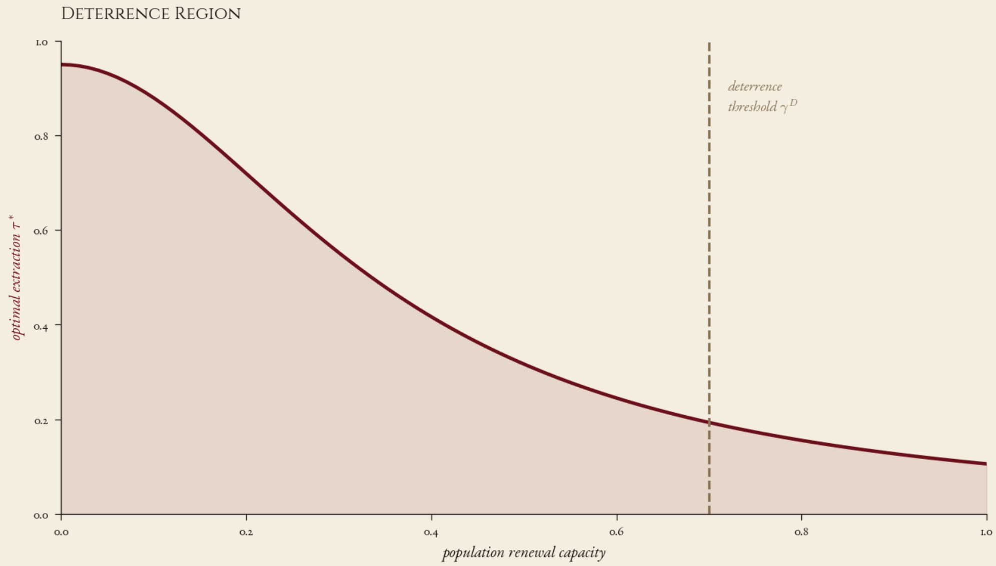

FIGURE 4.5. *Module 5: Deterrence analysis of renewal threat. Substrate behavior is constrained by anticipated renewal threat only when renewal capacity is in the deterrence region, with the deterrence operating dimension-specifically.*

The substantive implication is that substrates may behave provocatively in dimensions where their authority is already high, while behaving more cautiously in dimensions where their authority is weaker. The framework's apparatus predicts dimensional asymmetry in substrate behavior: substrates take risks in their strong dimensions and play it safe in their weak dimensions.

4.6.2 Empirical Patterns of Provocation

The historical record exhibits patterns consistent with this prediction.

The American substrate's behavior on regulatory scope (where substrate authority is strong) has been more provocative than its behavior on direct fiscal extraction (where substrate authority is weaker due to electoral discipline on tax increases). Regulatory expansion frequently faces opposition that the substrate then defeats, producing further regulatory consolidation. Fiscal expansion faces more directly defeating opposition through electoral consequence, and substrate behaves more cautiously on fiscal matters as a result.

The European Union substrate's behavior on supranational regulatory authority (where substrate authority is strong) has been more provocative than its behavior on direct member-state competence (where substrate authority is weaker due to treaty constraints). Regulatory expansion produces resistance that the substrate then defeats through European Court of Justice decisions, producing further consolidation. Member-state competence challenges face more directly defeating opposition through national parliamentary processes.

The Russian substrate's behavior on territorial-control authority (where substrate authority is strong) has been more provocative than its behavior on economic management authority (where substrate authority is weaker due to its dependence on global markets). Territorial expansion provokes resistance that the substrate then defeats militarily, producing further territorial consolidation.

Economic management faces less directly defeating opposition but operates within constraints that substrate cannot easily contest.

4.6.3 The Asymmetry Result

PROPOSITION 4.2 *(Dimensional Asymmetry of Substrate Behavior). Substrate behavior is more provocative in dimensions where substrate authority is high and the constraint apparatus's response capacity is low. Substrate behavior is more cautious in dimensions where substrate authority is low or the constraint apparatus's response capacity is high.*

The proposition follows from substrate's optimization. The expected gain from provocation depends on defeat probability (high A_d) and on consolidation magnitude (low constraint apparatus response capacity). Substrate's optimal behavior is therefore more provocative in dimensions where both conditions are met.

The substantive consequence is that substrate behavior is not uniform across its variables; it varies systematically with the substrate's strength and the constraint apparatus's configuration. The framework's prescriptive program (Chapter 5) is therefore concerned with maintaining constraint apparatus response capacity in each behavioral channel, not on extraction alone.

Part Two: Module 6 (Dimensional Renewal Capacity and Threshold Inclusion)

4.7 Dimensional Renewal Capacity

Module 5 established that renewal events have dimensional structure: renewal attempts target specific substrate dimensions, with the targeting affecting both the renewal demand and the substrate's response strategy. This module characterizes the population's renewal capacity in dimensional terms, treating renewal capacity as a multidimensional vector indexed by substrate dimension.

4.7.1 The Renewal Capacity Matrix

Renewal capacity is a multidimensional vector $\mathbf{C} = (C_1, C_2, \ldots, C_m)$ indexed by legitimacy resources, with each component measuring the population's capacity along one specific resource axis. The framework's apparatus adds dimensional structure on top of this, distinguishing renewal capacity for different substrate dimensions.

For each substrate dimension $d \in \{\tau, S, Q\}$, define the dimensional renewal capacity vector

$$\mathbf{C}_d = (C_{d,1}, C_{d,2}, \ldots, C_{d,m}),$$

where $C_{d,j}$ measures the population's renewal capacity along legitimacy resource j when the renewal target is dimension d. The full renewal capacity structure is therefore a $3 \times m$ matrix indexed by both substrate dimension and legitimacy resource:

$$\mathbf{C} = \begin{pmatrix} C_{\tau,1} & C_{\tau,2} & \cdots & C_{\tau,m} \\ C_{S,1} & C_{S,2} & \cdots & C_{S,m} \\ C_{Q,1} & C_{Q,2} & \cdots & C_{Q,m} \end{pmatrix}.$$

Each row characterizes the population's renewal capacity for renewal attempts targeting one substrate dimension. The rows are not independent: legitimacy resources contribute to renewal capacity for whichever dimensions they primarily support, with the contributions characterized through the resource-to-parameter mappings of Module 3.

4.7.2 Resource Contributions to Renewal Capacity

The legitimacy resources contribute to renewal capacity in each dimension according to their dimensional roles established in Module 3.

Founding mythology contributes primarily to scope-targeting renewal capacity ($C_{S,1}$). The shared narrative about substrate's legitimate purposes provides the substantive standard against which scope creep is measured, and renewal attempts targeting scope draw heavily on founding mythology for their legitimacy claims. Founding mythology contributes secondarily to extraction-targeting renewal capacity ($C_{\tau,1}$, where founding stories include implicit limits on extraction) and to quality-targeting renewal capacity ($C_{Q,1}$, where founding stories include expectations of adequate performance).

Procedural authority contributes to renewal capacity for each substrate variable roughly equally. Judicial review can support extraction-targeting renewal (statutory limits on what substrate can demand), scope-targeting renewal (constitutional limits on substrate authority), and quality-targeting renewal (statutory requirements for adequate performance). Legislative oversight similarly supports renewal capacity in each channel. Procedural authority is the most multifunctional resource for renewal capacity.

Expertise authority contributes primarily to quality-targeting renewal capacity ($C_{Q,3}$). Professional bodies provide the technical evaluation needed to mount renewal attempts targeting quality failures in their domains. Medical professional bodies can document inadequate health policy performance. Engineering societies can document infrastructure policy failures. Expertise authority contributes secondarily to scope-targeting renewal (professional views on legitimate substrate scope in their domains) and to extraction-targeting renewal (professional standards for appropriate compensation).

Democratic authority contributes to renewal capacity in each behavioral channel through population aggregation. Voting mechanisms allow population to express judgments about substrate performance on extraction, scope, or quality, with electoral consequence operating as a general-purpose renewal channel. Civil society organizations aggregate population concerns about specific substrate behavior on each variable.

Traditional authority contributes primarily to scope-targeting renewal capacity ($C_{S,5}$) through cultural-historical patterns articulating what substrate behavior is appropriate. Religious institutions and cultural traditions provide the deep value substrate against which scope creep is evaluated. Traditional authority contributes secondarily to other dimensions.

Crisis authority contributes to renewal capacity in each behavioral channel in extraordinary circumstances. Investigative journalism, special prosecutors, congressional hearings, and constitutional convention mechanisms can mobilize renewal capacity across the full spectrum when standard mechanisms fail. Crisis authority is the surge capacity for renewal in any channel.

4.7.3 Aggregating Resources to Renewal Capacity

The dimensional renewal capacity for each dimension aggregates the contributions from individual legitimacy resources:

$$\Phi_d = g_d(C_{d,1}, C_{d,2}, \ldots, C_{d,m}),$$

where g_d is the aggregation function for dimension d. The aggregation function depends on the institutional configuration and on how the resources interact in producing collective renewal capacity.

The aggregation is generally not additive: renewal capacity in any dimension depends on having multiple resources contributing, with the resources operating in different modalities. A polity with high $C_{S,1}$ (founding mythology) but low $C_{S,2}$ (procedural authority for scope) and low $C_{S,5}$ (traditional authority) has low overall Φ_S because scope-targeting renewal requires multiple modal sources of legitimacy. The aggregation function reflects this multimodal requirement.

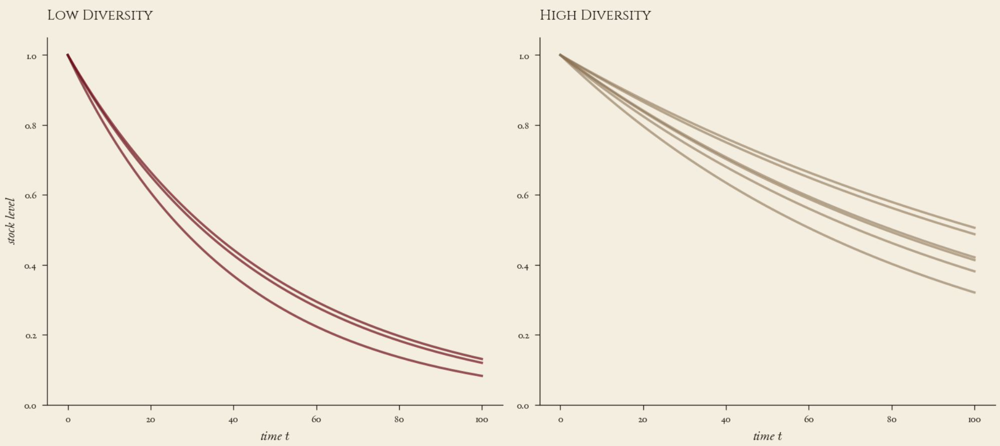

FIGURE 4.6. *Module 3: Resource diversity dynamics. Higher diversity produces more cascade-resistant resource portfolios, with the diversity required for cascade resistance characterized through the coverage graph topology developed in Module 4.*

4.8 Thresholds Across Dimensions

The framework introduces three threshold specifications on renewal capacity: feasibility, probable-success, and deterrence. The thresholds operate separately for each substrate dimension.

4.8.1 Feasibility Threshold

The feasibility threshold for dimension d is the renewal capacity level Φ_d^F above which mounting a renewal attempt targeting dimension d is technically possible. Below this threshold, the population cannot organize a coherent renewal attempt: insufficient legitimacy resources, insufficient coordination capacity, insufficient mobilization potential. The renewal attempt cannot be launched in any meaningful form.

The feasibility region for dimension d is

$$\mathcal{F}_d = \{\mathbf{\Phi} : \Phi_d \geq \Phi_d^F\}.$$

A polity is in $\mathcal{F}_d$ when its renewal capacity for dimension d exceeds the feasibility threshold. The polity may not be able to launch successful renewal attempts targeting d, but it can at least mount such attempts.

Empirical instances of $\mathcal{F}_d$. Polities in the feasibility region for some dimensions but not others are common. A polity with strong democratic authority and crisis authority has Φ_τ above feasibility (extraction-targeting renewal can be launched through electoral and crisis channels) but may have Φ_S below feasibility if founding mythology and traditional authority have substantially decayed. The polity can mount tax revolts but cannot mount constitutional renewal of legitimate scope.

4.8.2 *Probable-Success Threshold*

The probable-success threshold for dimension d is the renewal capacity level Φ_d^P above which a mounted renewal attempt targeting dimension d has greater than 50% probability of achieving its renewal demand. This threshold is more demanding than feasibility because mounting an attempt is not the same as winning it. The substrate's defeat capacity (Module 5) operates against the renewal attempt, with defeat probability depending on substrate authority and population renewal capacity.

The probable-success region for dimension d is

$$\mathcal{P}_d = \{\boldsymbol{\Phi} : P(\text{success}_d|\text{attempt}_d) > 0.5\},$$

where the probability depends on Φ_d, substrate authority A_d, and the substrate's optimal response. The probable-success threshold Φ_d^P depends on A_d: higher substrate authority on dimension d raises the threshold required for population to have probable-success capacity.

Empirical instances of $\mathcal{P}_d$. Polities in the probable-success region for some dimensions but not others exhibit the dimensional asymmetry of substrate's defeat capacity. A polity with weak A_τ but strong A_S may be in $\mathcal{P}_\tau$ (extraction-targeting renewal is likely to succeed) but not in $\mathcal{P}_S$ (scope-targeting renewal is likely to fail because substrate has high defeat capacity in scope). The polity's effective renewal capacity is dimension-specific.

4.8.3 *Deterrence Threshold*

The deterrence threshold for dimension d is the renewal capacity level Φ_d^D above which substrate's optimal behavior on dimension d is constrained by anticipated renewal threat. This is the most demanding threshold because deterrence requires substrate to perceive the renewal threat as sufficient to discipline its behavior even when no renewal attempt is currently being mounted.

The deterrence region for dimension d is

$$\mathcal{D}_d = \{\boldsymbol{\Phi} : \text{substrate's optimal behavior on dimension } d \text{ is constrained by renewal threat}\}.$$

Whether the polity is in $\mathcal{D}_d$ depends not only on the renewal capacity Φ_d but also on substrate's perception of the threat (which depends on substrate's information about renewal capacity), substrate's discount factor (which determines how heavily substrate weights anticipated threats), and substrate's optimization parameters (which determine how much constraint the threat produces).

Empirical instances of $\mathcal{D}_d$. Polities in the deterrence region for any variable are rare in advanced operational phase. Most contemporary polities have substrate operating at extraction, scope, and quality choices substantially above what deterrence would produce, indicating that renewal threat is not effectively constraining substrate behavior. The framework's prescriptive program targets restoration of deterrence-region renewal capacity in each behavioral channel.

4.8.4 *The Threshold Vector*

The full threshold structure for the polity is a vector of thresholds, one per substrate variable:

$$\mathbf{\Phi}^F = (\Phi_\tau^F, \Phi_S^F, \Phi_Q^F), \quad \mathbf{\Phi}^P = (\Phi_\tau^P, \Phi_S^P, \Phi_Q^P), \quad \mathbf{\Phi}^D = (\Phi_\tau^D, \Phi_S^D, \Phi_Q^D).$$

A polity's overall threshold profile is its location relative to these three vectors. The polity may be in different threshold regions for different substrate variables, producing the verdict asymmetry that the next section develops.

4.9 Strict Inclusion Within Each Dimension

The strict inclusion result $\mathcal{D} \subset \mathcal{P} \subset \mathcal{F}$ holds within each substrate dimension separately.

THEOREM 4.3 (Dimensional Threshold Inclusion). *For each substrate dimension* $d \in \{\tau, S, Q\}$ *and generic parameter values:*

$$\mathcal{D}_d \subset \mathcal{P}_d \subset \mathcal{F}_d,$$

with strict inclusion. The deterrence region for dimension d *is properly contained in the probable-success region, which is properly contained in the feasibility region.*

Proof. We prove $\mathcal{P}_d \subset \mathcal{F}_d$ first, then $\mathcal{D}_d \subset \mathcal{P}_d$, with strictness in each case.

Proof that $\mathcal{P}_d \subseteq \mathcal{F}_d$. If the polity is in $\mathcal{P}_d$, then a renewal attempt targeting dimension d has greater than 50% probability of success. For success to be possible, the renewal attempt must be mountable, which requires $\mathbf{\Phi} \in \mathcal{F}_d$. Therefore $\mathcal{P}_d \subseteq \mathcal{F}_d$.

Proof of strict inclusion $\mathcal{P}_d \subsetneq \mathcal{F}_d$. Consider a polity with Φ_d at the feasibility threshold Φ_d^F but with A_d at a sufficiently high level. The substrate's defeat capacity p_{d,r^*} is high enough that even a feasible renewal attempt has less than 50% probability of success. Such a polity is in $\mathcal{F}_d$ but not in $\mathcal{P}_d$. Strictness follows.

Proof that $\mathcal{D}_d \subseteq \mathcal{P}_d$. If the polity is in $\mathcal{D}_d$, then substrate's optimal behavior on dimension d is constrained by anticipated renewal threat. For substrate to be constrained, substrate must perceive the renewal threat as substantively meaningful, which requires the threat to be sufficient that substrate's expected post-renewal authority on dimension d would fall if substrate behaved more aggressively. Substrate's expected post-renewal authority depends on probable success of renewal attempts that would be triggered by aggressive substrate behavior. For the threat to be substantively meaningful, probable success must be greater than 50%. Therefore $\mathbf{\Phi} \in \mathcal{P}_d$. Therefore $\mathcal{D}_d \subseteq \mathcal{P}_d$.

Proof of strict inclusion $\mathcal{D}_d \subsetneq \mathcal{P}_d$. Consider a polity with renewal capacity Φ_d exceeding the probable-success threshold Φ_d^P but with substrate having high time preference (low δ_S) or with substrate not perceiving the threat (asymmetric information). Substrate's optimal behavior is not effectively constrained by the renewal threat even though probable success would obtain. Such a polity is in $\mathcal{P}_d$ but not in $\mathcal{D}_d$. Strictness follows.

Combining the two inclusions and strictness:

$$\mathcal{D}_d \subsetneq \mathcal{P}_d \subsetneq \mathcal{F}_d.$$

■

The proof applies within each substrate dimension separately. The substantive interpretation is that the threshold structure operates within each dimension: a polity has different threshold profiles for different dimensions, with strict inclusion holding within each dimension's profile.

4.9.1 *The Threshold Profile*

The dimensional threshold inclusion characterizes the polity by three dimensional threshold profiles with potentially different locations rather than by a single threshold profile.

Best-case profile. A polity with renewal capacity in $\mathcal{D}_d$ for each $d \in \{\tau, S, Q\}$ has effective deterrence across substrate behavior. Substrate's optimal behavior is constrained by renewal threat in each channel. This is what the framework's prescriptive program aims for, and it is the configuration that produces the long-run trajectory near the Irreducible Floor on each variable.

Asymmetric dimensional profile. A polity in $\mathcal{D}_\tau$ but only $\mathcal{P}_S$ and $\mathcal{F}_Q$ has effective deterrence on extraction but only probable-success capacity on scope and only feasibility capacity on quality. Substrate's behavior is constrained on extraction but less constrained on scope (probable success but not deterrence) and minimally constrained on quality (only feasibility). This is the typical pattern of polities where electoral consequences discipline tax extraction but where scope creep proceeds with less effective constraint and where quality failure goes largely undetected and unaddressed.

Worst-case dimensional profile. A polity outside $\mathcal{F}_d$ for any dimension cannot mount renewal attempts targeting that dimension at all. Substrate's behavior in that dimension is unconstrained by renewal threat. This is the configuration where substrate captures dimension d entirely, with no population response capacity remaining.

The dimensional threshold inclusion thus produces a 3 × 3 grid of possible threshold positions (each dimension can be in $\mathcal{D}_d$, $\mathcal{P}_d \setminus \mathcal{D}_d$, $\mathcal{F}_d \setminus \mathcal{P}_d$, or outside $\mathcal{F}_d$), allowing 27 distinct threshold profiles for the polity. The framework's apparatus characterizes which profiles are stable, which produce specific trajectories, and which are most vulnerable to terminal capture.

4.10 Dimensional Verdict Asymmetry

Different observers using different threshold specifications can reach different verdicts about the same polity. Verdict asymmetry is enriched through dimensional variation: observers tracking different dimensions can systematically disagree about the same polity.

4.10.1 *Verdict Asymmetry Through Threshold Inclusion*

Verdict asymmetry through threshold inclusion: an observer using the deterrence threshold says the polity is in trouble whenever $\mathbf{\Phi} \notin \mathcal{D}$. An observer using the probable-success threshold says the polity is in trouble whenever $\mathbf{\Phi} \notin \mathcal{P}$. An observer using the feasibility threshold says the polity is in trouble whenever $\mathbf{\Phi} \notin \mathcal{F}$. The three observers' verdicts differ by inclusion: more polities are in trouble under the deterrence verdict than under the probable-success verdict, and more under probable-success than under feasibility.

4.10.2 *Observer Dimensional Preference*

Dimensional verdict asymmetry adds observer dimensional preferences to the threshold-based asymmetry. An observer can prioritize one dimension over others when forming their overall verdict. Different observers tracking different dimensions can reach systematically different overall verdicts about the same polity.

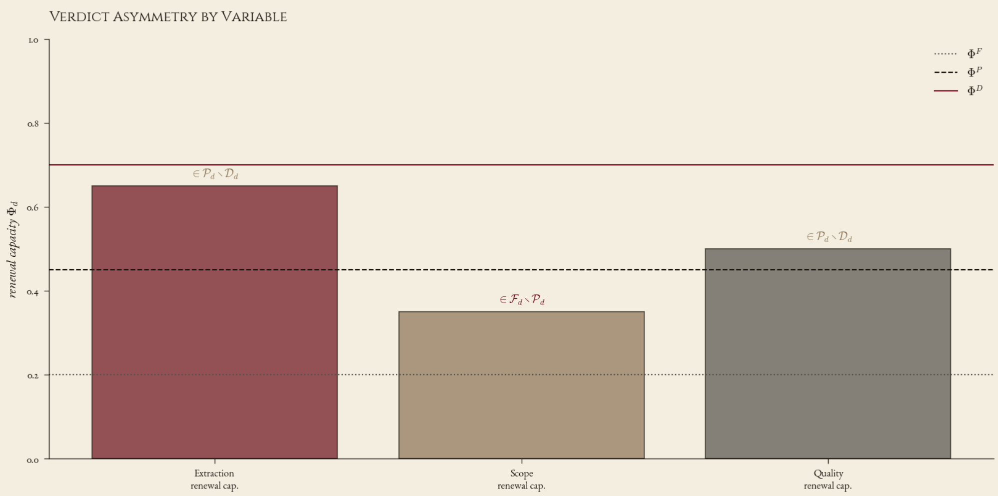

FIGURE 4.7. *Module 6: Dimensional verdict asymmetry. The same polity can be evaluated differently by observers prioritizing different substrate dimensions. The libertarian-progressive disagreement about American polity health, the institutional-populist disagreement about contemporary democracies, and the realist-idealist disagreement about international polities are all explained through different dimensional emphases rather than factual disagreement.*

Extraction-focused observer. An observer concerned primarily with substrate's rent extraction evaluates the polity primarily by $\Phi_{\mathcal{T}}$. The observer considers the polity in trouble when $\mathbf{\Phi} \notin \mathcal{D}_{\mathcal{T}}$, $\mathcal{P}_{\mathcal{T}}$, or $\mathcal{F}_{\mathcal{T}}$ depending on threshold preference.

Scope-focused observer. An observer concerned primarily with substrate's scope creep evaluates the polity primarily by Φ_S. The observer considers the polity in trouble when $\mathbf{\Phi} \notin \mathcal{D}_S, \mathcal{P}_S$, or $\mathcal{F}_S$ depending on threshold preference.

Quality-focused observer. An observer concerned primarily with substrate's quality of legitimate function performance evaluates the polity primarily by Φ_Q. The observer considers the polity in trouble when $\mathbf{\Phi} \notin \mathcal{D}_Q, \mathcal{P}_Q$, or $\mathcal{F}_Q$ depending on threshold preference.

The empirically observed pattern of partisan disagreement about polity health is partially explained by the dimensional verdict asymmetry. Different political traditions emphasize different substrate dimensions, leading to systematic differences in their evaluation of the same polity.

4.10.3 *Empirical Patterns of Verdict Asymmetry*

Several empirically observed disagreements about polity health are explained by the dimensional verdict asymmetry.

The libertarian–progressive disagreement about American polity health. Libertarian observers tend to emphasize scope (substrate's expansion beyond legitimate rights-protective functions) and extraction (substrate's resource consumption). Progressive observers tend to emphasize quality (substrate's performance of legitimate functions, particularly those they consider legitimate). Observing the same American polity, libertarians often see substantial trouble (substrate extraction is high, scope is far beyond legitimate) while progressives see less trouble (substrate quality is mixed but improving in some domains). The disagreement is partially explained by different dimensional emphasis rather than by disagreement about facts.

The institutional versus populist disagreement about contemporary democracies. Institutional observers emphasize procedural authority and quality of legitimate function performance. Populist observers emphasize democratic authority and sometimes traditional authority. Observing the same polity, institutional observers see substantial constraint capacity (procedural authority is functioning, quality is adequate in many domains) while populist observers see substantial trouble (democratic authority is captured by elites, traditional authority is being eroded). The disagreement reflects different dimensional and resource emphases.

The realist versus idealist disagreement about international polities. Realist observers emphasize substrate's effectiveness at core security functions. Idealist observers emphasize substrate's adherence to legitimate scope and rights protection. Observing the same polity, realists may see effective substrate (high quality on security, effective on extraction) while idealists may see compromised substrate (scope creep into illegitimate domains, extraction beyond what rights-protective functions require). The disagreement reflects different dimensional emphases applied to international comparison.

4.10.4 *The Verdict Disagreement Result*

PROPOSITION 4.4 (Dimensional Verdict Disagreement). *Two observers A and B tracking the same polity can reach systematically different overall verdicts about the polity if they prioritize different substrate dimensions in their evaluation. The disagreement is not about facts but about which dimension is most important for evaluating polity health.*

The proposition combines two sources of verdict disagreement: threshold preference (deterrence vs probable-success vs feasibility) and dimensional preference (which substrate dimension matters most). Both contribute to why observers reach different verdicts about the same polity.

4.10.5 The Implications for Discourse

The dimensional verdict asymmetry has implications for political discourse about polity health. Disagreements about whether a polity is healthy or not are typically framed as factual disagreements (about how much substrate is extracting, expanding, or failing). The framework's analysis suggests these are often value disagreements about which dimension matters most.

The framework's prescriptive program (Chapter 5) addresses this by treating each substrate variable as analytically important. The framework's verdict on polity health depends on the polity's full threshold profile, not on any single variable. A polity in trouble on any one variable is at risk of trajectory toward terminal capture on that variable, regardless of how it performs on the others.

The framework's analytical position is therefore that observers focused on any single substrate variable are missing important features of polity health. A complete evaluation requires assessing the polity across the full threshold profile rather than any single variable's location.

4.11 Renewal Capacity Decay Trajectories

Renewal capacity decays over operational time as legitimacy resources deplete. The framework's apparatus characterizes decay trajectories in each behavioral channel, which can proceed at different rates depending on which resources are depleting fastest.

4.11.1 Decay Trajectory Within a Dimension

For each substrate dimension d, the renewal capacity Φ_d decays over operational time according to the depletion of the legitimacy resources contributing to dimension d's renewal capacity. The decay rate depends on:

Rate of legitimacy resource capture. Substrate's optimal capture rates from Module 3 determine how fast each resource depletes. Resources whose primary contribution is to dimension d depletion produces Φ_d falling at rates proportional to how fast those resources are captured.

Refresh capacity within each dimension. Population's refresh capacity for resources supporting dimension d determines how much of the depletion is offset. Dimensions whose supporting resources receive less refresh effort (per Module 3's allocation analysis) experience faster Φ_d decay.

Cascade dynamics within each dimensional subgraph. Module 4's cascade vulnerability analysis applies to each dimensional subgraph G_d. Cascade dynamics in G_d produce accelerating decay in Φ_d as cascading captures undermine detection capacity for dimension d.

4.11.2 Threshold Crossings as Trajectory Markers

As Φ_d decays, the polity crosses threshold boundaries in dimension d, moving from $\mathcal{D}_d$ to $\mathcal{P}_d$ to $\mathcal{F}_d$ to outside $\mathcal{F}_d$. Each crossing marks a substantive change in substrate's optimization on dimension d.

FIGURE 4.8. *Module 6: Renewal capacity decay trajectory and threshold crossings. As legitimacy resources deplete, the polity crosses deterrence, probable-success, and feasibility thresholds, producing substantive changes in substrate's optimization at each crossing.*

Crossing Φ_d^D from above. Substrate's optimal behavior on dimension d is no longer constrained by anticipated renewal threat. Substrate becomes more aggressive on dimension d. The substrate's optimization shifts toward higher τ, larger S, or lower Q depending on which dimension. The polity transitions from constrained substrate behavior to unconstrained-by-deterrence behavior.

Crossing Φ_d^P from above. Renewal attempts targeting dimension d are no longer probably-successful. Substrate's defeat capacity has grown sufficient that mounted renewal attempts will likely fail, producing renewal capture (Module 5) rather than reset. The polity transitions from probably-successful renewal capacity to renewal capture probability.

Crossing Φ_d^F from above. Renewal attempts targeting dimension d are no longer feasible. Population cannot organize coherent renewal attempts on dimension d. The polity transitions to substrate behavior on dimension d being unconstrained by renewal possibility entirely.

The polity's trajectory across these crossings characterizes its decay path in each dimension. Different dimensions can be at different points in this trajectory, producing the dimensional asymmetric profiles characterized in the previous section.

4.11.3 The Compound Trajectory

The polity's overall renewal capacity trajectory is a composite of the dimensional trajectories. As the dimensions cross their respective thresholds at different rates, the polity moves through a sequence of dimensional asymmetric profiles, with the final terminal capture configuration occurring when all dimensions have crossed below feasibility.

The framework's apparatus characterizes the typical trajectory under partial implementation regimes. Without renewal capacity refresh (Module 7's Prescription 3), the cascade dynamics produce progressive decay across dimensions, with the dimension whose supporting resources are most cascade-vulnerable falling first. Subsequent decay in other dimensions follows according to the cascade pattern in their respective subgraphs.

The empirical pattern of polities undergoing dimensional decay at different rates (Roman Republic's rapid scope-detection cascade, Han's rapid quality-detection cascade, contemporary American mixed pattern) is consistent with this prediction. Different polities exhibit different dimensional trajectories depending on which resources are most cascade-vulnerable in their specific institutional configurations.

4.12 Renewal Capacity and the Prescriptive Program

The renewal capacity analysis connects directly to the framework's prescriptive program (Chapter 5) through Prescription 3 (refresh capacity), which must operate in each behavioral channel.

4.12.1 Refresh Capacity

Refresh capacity is the institutional infrastructure for maintaining legitimacy resources over operational time. Refresh must operate on resources supporting each behavioral channel.

Refresh for extraction-detecting resources. Procedural authority resources that support extraction detection (judicial review, legislative oversight) require institutional support including selection mechanisms, jurisdiction protection, resource allocation, and procedural protection from substrate interference. Refresh institutions for these resources include constitutional protections for judicial independence, legislative procedure rules, transparency mechanisms, and electoral monitoring.

Refresh for scope-detecting resources. Founding mythology and traditional authority resources that support scope detection require institutional support including educational institutions transmitting the founding narrative, religious freedom protecting traditional value sources, legal traditions sustaining articulation of legitimate scope, and historical memory institutions preserving lessons. Refresh institutions for these resources include constitutional rights protecting religious and educational institutions from substrate capture, robust civil society organization, historical memory institutions like archives and historical societies.

Refresh for quality-detecting resources. Expertise authority resources that support quality detection require institutional support including professional bodies with formal continuity, peer review mechanisms, professional standards enforcement, technical expertise transmission, and protection from substrate interference. Refresh institutions for these resources include guild-style professional self-governance, accreditation independence, scientific publication freedom, and protection of academic institutions from political interference.

4.12.2 Cross-Dimensional Refresh Coordination

Refresh effort requires coordination because legitimacy resources contribute to multiple behavioral channels through their detection roles (Module 3). Refresh effort allocated to procedural authority supports renewal capacity in each channel; refresh effort allocated to expertise authority primarily supports the quality channel. Refresh allocation must consider balance across channels to avoid over-investing in one while under-investing in others.

The framework's prescriptive program recommends refresh allocation that maintains dimensional balance, with attention to which dimensions are most cascade-vulnerable in the polity's specific configuration. Polities with cascade-vulnerable scope detection (low G_S resilience) need to direct refresh effort toward founding mythology, traditional authority, and scope-supporting procedural authority. Polities with cascade-vulnerable quality detection (low G_Q resilience) need to direct refresh effort toward expertise authority and quality-supporting procedural authority.

CHAPTER V

Comparative Statics, Prescriptions, and Implementation

This chapter derives the framework's prescriptive program from the comparative-statics analysis of the previous chapters. The chapter combines Module 7's comparative-statics analysis with the institutional-design directive that addresses the constraint-quality trade-off.

The chapter's analytical contributions are several. Each of the four comparative-statics prescriptions (resource diversity, capture detection, refresh capacity, capture rotation) is shown to produce favorable effects on extraction, scope, and quality through different specific mechanisms. Parallel constraint design is added to the prescriptive program, addressing the procedural friction and substantive resource competition channels through which the constraint apparatus would otherwise reduce substrate quality. The regime simulations produce outcome vectors (τ^*, S^*, Q^*) that characterize each regime's predicted substrate behavior.

5.1 The Four Prescriptions

Four prescriptions follow from comparative statics on substrate's optimization. Each prescription's effects are derived for each substrate variable, characterizing how it affects extraction, scope, and quality.

5.1.1 Prescription 1: Resource Diversity

The first prescription is resource diversity: the constraint apparatus should consist of multiple legitimacy resources operating in distinct modalities, with the resources collectively supporting detection across all substrate dimensions.

Effect on extraction. Higher diversity of constraint resources targeting extraction detection (procedural authority, expertise authority, democratic authority, crisis authority) raises detection capacity for extraction, which through Module 2's first-order condition lowers optimal extraction. The comparative-static result is $\partial\tau^*/\partial J_\tau < 0$, where J_τ is diversity in the extraction-detecting subgraph G_τ.

Effect on scope. Higher diversity of constraint resources targeting scope detection (founding mythology, traditional authority, procedural authority, crisis authority) raises detection capacity for scope creep, which raises the marginal cost to substrate of scope expansion. The comparative-static result is $\partial S^*/\partial J_S < 0$, where J_S is diversity in the scope-detecting subgraph G_S. Substrate enters fewer illegitimate domains when scope-detection capacity is higher.

Effect on quality. Higher diversity of constraint resources targeting quality detection (expertise authority, procedural authority, democratic authority, crisis authority) raises detection capacity for quality failure, which raises the marginal cost to substrate of quality underinvestment. The comparative-static result is $\partial Q^*/\partial J_Q > 0$, where J_Q is diversity in the quality-detecting subgraph G_Q. Substrate invests more in quality when quality-detection capacity is higher.

Prescription 1 specifies that diversity must be maintained within each dimensional subgraph. Diversity in the full coverage graph that does not extend to each dimensional subgraph produces dimension-specific failures even when overall diversity appears adequate.

5.1.2 Prescription 2: Capture Detection

The second prescription is capture detection: the constraint apparatus should maintain low opacity, supporting transparent observation of substrate behavior across all dimensions.

Effect on extraction. Lower opacity raises extraction-detection capacity, which lowers optimal extraction. The comparative-static result is $\partial\tau^*/\partial v < 0$, where v is opacity.

Effect on scope. Lower opacity raises scope-detection capacity, which lowers optimal scope expansion. Scope creep is harder to hide when opacity is low, raising substrate's marginal cost of entering illegitimate domains. The comparative-static result is $\partial S^*/\partial v < 0$.

Effect on quality. Lower opacity raises quality-detection capacity, which raises optimal quality investment. Quality failure is harder to hide when opacity is low, raising substrate's marginal cost of underinvestment. The comparative-static result is $\partial Q^*/\partial v > 0$.

Capture detection therefore has favorable effects on each variable through a common mechanism: lower opacity raises detection capacity uniformly across whichever channels depend on observation of substrate behavior.

5.1.3 Prescription 3: Refresh Capacity

The third prescription is refresh capacity: the constraint apparatus should be sustained over operational time through deliberate institutional support for legitimacy resources, preventing the cascading depletion that otherwise occurs.

Effect on extraction. Higher refresh capacity for extraction-supporting resources (procedural authority, democratic authority, crisis authority) sustains extraction-detection capacity over operational time. Without refresh, these resources deplete through the Module 3 cascade dynamics, producing rising τ^*. The comparative-static result is $\partial\tau^*_{long-run}/\partial\mu_\tau < 0$, where μ_τ is refresh capacity for extraction-supporting resources.

Effect on scope. Higher refresh capacity for scope-supporting resources (founding mythology, traditional authority, scope-related procedural authority) sustains scope-detection capacity over operational time. Without refresh, these resources deplete through cascade dynamics, producing rising S^*. The comparative-static result is $\partial S^*_{long-run}/\partial \mu_S < 0$.

Effect on quality. Higher refresh capacity for quality-supporting resources (expertise authority, quality-related procedural authority) sustains quality-detection capacity over operational time. Without refresh, these resources deplete, producing falling Q^*. The comparative-static result is $\partial Q^*_{long-run}/\partial \mu_Q > 0$.

Refresh capacity therefore has favorable effects on each variable, with the magnitude depending on which channels' supporting resources receive refresh effort. The framework's prescriptive program requires refresh effort to be allocated across channels, with attention to which channels are most cascade-vulnerable in the polity's specific configuration.

5.1.4 Prescription 4: Capture Rotation

The fourth prescription is capture rotation: the constraint apparatus should include mechanisms that prevent persistent capture of constraint resources, with rotation operating through term limits, sortition, procedural reset mechanisms, and federalism-enabled mobility.

Effect on extraction. Higher rotation reduces persistent capture of extraction-detecting resources, which reduces the cumulative extraction-favorable bias of the constraint apparatus. The comparative-static result is $\partial \tau^*_{long-run}/\partial \eta_\tau < 0$, where η_τ is rotation in extraction-supporting resources.

Effect on scope. Higher rotation reduces persistent capture of scope-detecting resources, which reduces the cumulative scope-favorable bias. The comparative-static result is $\partial S^*_{long-run}/\partial \eta_S < 0$.

Effect on quality. Higher rotation reduces persistent capture of quality-detecting resources, which reduces the cumulative quality-failure-favorable bias. The comparative-static result is $\partial Q^*_{long-run}/\partial \eta_Q > 0$.

Capture rotation has favorable effects on each variable, again with the magnitude depending on which channels' supporting resources are rotated.

5.1.5 Summary

The four prescriptions all have favorable effects on each substrate variable through their effects on detection capacity. Each prescription operates through a different mechanism (resource portfolio, opacity, refresh, rotation), and each produces comparative-statics improvements in substrate behavior on each variable through that mechanism.

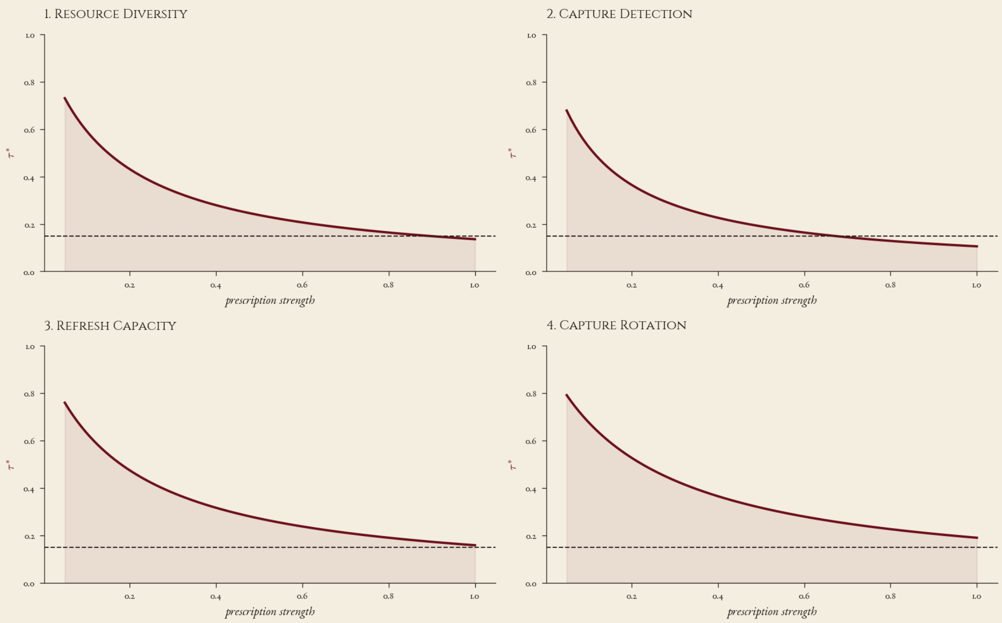

FIGURE 5.1. *Module 7: The four comparative-statics prescriptions. Resource diversity, capture detection, refresh capacity, and capture rotation each produce favorable comparative-statics outcomes through different mechanisms in the constraint apparatus.*

Each prescription's effect on each substrate variable depends on whether the prescription's specific implementation addresses that variable's supporting resources. A polity implementing diversity only for extraction-detecting resources produces favorable extraction outcomes but does not improve scope or quality outcomes. A polity implementing refresh only for procedural authority produces favorable outcomes on each variable (since procedural authority supports detection in each channel) but with stronger effects where procedural authority dominates the supporting resource portfolio.

5.2 Parallel Constraint Design

The four comparative-statics prescriptions characterize what the constraint apparatus must do. A separate question addresses how the constraint apparatus must be designed to actually do those things.

5.2.1 The Constraint-Quality Trade-off

The constraint apparatus affects substrate quality through two channels identified in earlier work. The procedural friction channel: constraint apparatus operations that require substrate decisions to undergo approvals, reviews, and consultations slow substrate's capacity for action, including action in legitimate domains. The substantive resource competition channel: constraint apparatus operations consume polity resources that would otherwise be available for legitimate substrate functions.

Both channels are real but neither is unconditional. The trade-off between constraint capacity and substrate quality can be managed through institutional design choices that affect each channel.

The procedural friction channel can be managed through institutional designs that operate constraint apparatus in parallel with substrate action rather than in series. Bidirectional communication between constraint resources and substrate operations allows constraint to operate continuously rather than as bottleneck checkpoints. Pre-clearance procedures for routine substrate actions reduce friction without reducing constraint capacity. Specialized constraint resources for specific substrate domains reduce cross-domain procedural overhead.

The substantive resource competition channel can be managed through institutional designs that fund constraint apparatus through dedicated mechanisms separate from substrate's primary operations. Constitutional guarantees for constraint resource funding insulate the apparatus from substrate's resource-allocation decisions. Independent revenue streams (filing fees, dedicated taxes, foundation funding) provide constraint apparatus with resources that do not compete with substrate's legitimate functions.

5.2.2 *Design Rules*

The constraint apparatus should be designed to operate in parallel with substrate action and with funding independent of substrate's primary operations, minimizing both procedural friction and substantive resource competition.

Parallel-not-serial design. Constraint apparatus operations should run in parallel with substrate action rather than as sequential checkpoints that substrate must pass through. Judicial review should occur during and after substrate action rather than as ex ante approval. Legislative oversight should operate through ongoing committee work rather than as ex ante authorization for each substrate decision. Press evaluation should be continuous rather than episodic.

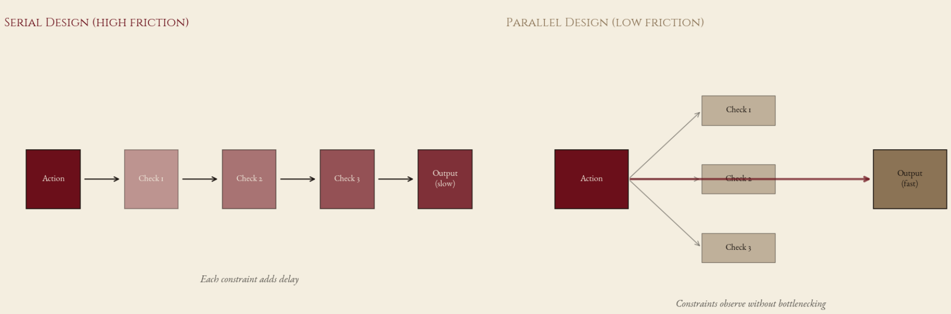

FIGURE 5.2. *Module 7: Parallel constraint design. Serial-then-sequential designs produce procedural friction that reduces substrate quality. Parallel-not-serial designs allow constraint apparatus to operate in parallel with substrate action, preserving constraint capacity without bottlenecking action.*

Independent funding. Constraint resources should have dedicated funding streams that do not compete with substrate's legitimate function budgets. Constitutional guarantees for judicial branch funding, independent press infrastructure (subscription-based or foundation-supported), professional bodies funded through membership and accreditation fees, civic society funded through philanthropic support all reduce the substantive resource competition channel.

Scaling with substrate scope. Constraint apparatus should scale with substrate scope, with new substrate domains paired with corresponding new constraint resources. The framework's analysis

(Module 4) shows that scope expansion without constraint expansion produces low effective dimensional diversity in the new domains. Scope expansion (where it occurs at all) should be paired with constraint apparatus expansion to maintain dimensional diversity.

5.2.3 *Venice, America, and the European Union*

Three cases illustrate parallel constraint design at work and at fault.

The Venetian Republic's deliberate institutional engineering produced a constraint apparatus that operated largely in parallel with substrate action. The Council of Ten could investigate without first obtaining authorization from the Doge or other authorities. The Senate could legislate without ex ante approval from the Council of Forty. The Inquisitors operated continuously rather than as response to specific events. The parallel design supported constraint capacity without substantial procedural friction, contributing to the Venetian Republic's extraordinary operational longevity.

The American post-1789 constitutional design has elements of both parallel and serial constraint structures. Judicial review operates largely in parallel (review of executive and legislative action ex post). Legislative oversight operates through both parallel mechanisms (ongoing committee work) and serial mechanisms (confirmation of executive appointments, budget appropriation). The mixed design has produced mixed results: parallel constraint mechanisms have remained more effective than serial ones over the operational phase.

The European Union's regulatory design has substantial procedural friction. Substrate action in many domains requires sequential approval through Commission, Council, and Parliament processes, with each stage providing constraint but also generating delay. The serial design produces high constraint capacity but at the cost of substrate quality in domains requiring rapid response. The Brexit decision and subsequent EU reform discussions have recognized this trade-off as a substantive concern about EU institutional design.

5.3 Cross-Dimensional and Cross-Prescription Interactions

The four prescriptions plus parallel constraint design interact across dimensions in ways that produce additional substantive predictions.

5.3.1 *Substitution and Compensation*

Module 2 established that substrate has multiple ways to acquire rent: higher τ within existing scope, larger S creating new scope, or lower Q at the cost of legitimate function performance. These three are partial substitutes in substrate's optimization. The framework's prescriptions must address each variable or substrate will compensate.

A polity that effectively constrains τ but not S produces substrate that compensates by expanding scope: substrate captures more rent through entering new domains rather than extracting more from existing domains. The empirical pattern of regulatory state expansion in polities with strong electoral discipline on tax extraction is consistent with this prediction.

A polity that effectively constrains S but not Q produces substrate that compensates by underinvesting in quality: substrate operates within legitimate scope but performs poorly, with the rent extraction occurring through reduced effort rather than through scope or extraction. The empirical pattern of substrate dysfunction in some constitutionally-constrained polities is consistent with this prediction.

A polity that effectively constrains τ and S but not Q produces substrate that compensates entirely through quality underinvestment. The framework's prescriptive program must address each substrate variable to prevent this compensation pattern.

5.3.2 *Reinforcement Across Prescriptions*

Joint implementation of multiple prescriptions produces effects that exceed the sum of individual implementations. Several specific reinforcement effects:

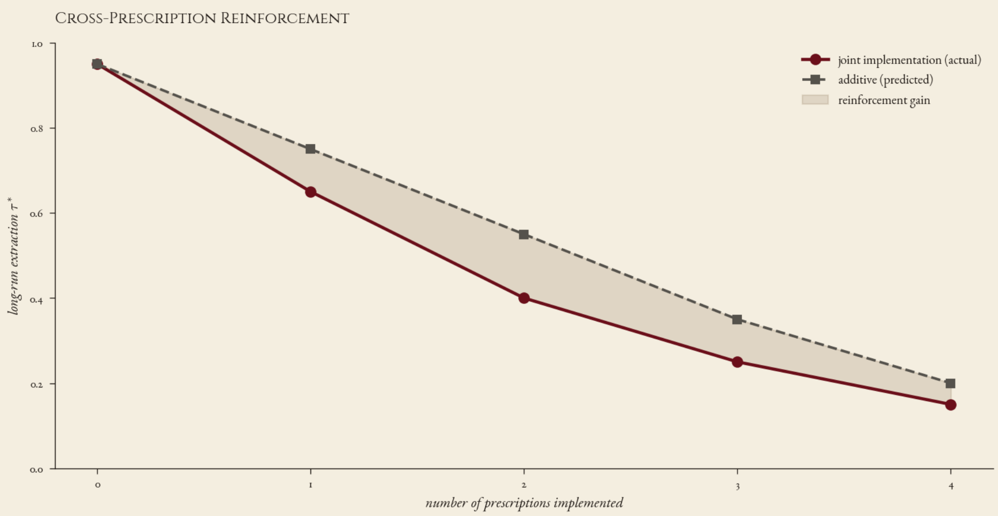

FIGURE 5.3. *Module 7: Cross-prescription interactions. Joint implementation produces reinforcement effects that exceed the sum of individual implementations, with specific interactions among the four prescriptions and parallel constraint design.*

Diversity and detection reinforcement. Diverse constraint resources produce more comprehensive detection of substrate behavior, which through the substrate's optimization produces more constrained substrate. Adding diversity strengthens detection's effects; adding detection makes diversity more impactful.

Refresh and rotation reinforcement. Refresh sustains the constraint apparatus over operational time; rotation prevents persistent capture within the apparatus. Joint implementation produces a constraint apparatus that is both sustained and uncaptured, which produces stronger constraint than either prescription alone.

Parallel design interactions. Parallel constraint design interacts with the four prescriptions by reducing the costs of implementation. Running constraint in parallel rather than serially with substrate action reduces the procedural friction associated with diverse constraint apparatus. Independent funding reduces the resource competition associated with sustained refresh. Parallel constraint design makes the four prescriptions more achievable in practice.

5.3.3 Trade-offs Among Prescriptions

Some trade-offs exist among the prescriptions and require institutional design choices. The most important is the trade-off between rotation and expertise.

Rotation prevents persistent capture but disrupts the institutional knowledge accumulation that produces high-quality constraint capacity. A judiciary with frequent rotation has lower individual judge expertise than a judiciary with longer tenure. A legislative committee with frequent membership rotation has lower oversight expertise than a committee with stable membership. The trade-off is real and the institutional design must balance rotation against expertise preservation.

The framework's resolution is that rotation should be applied selectively rather than uniformly. Procedural authority resources that operate primarily through general principles (constitutional review) can sustain higher rotation without large expertise loss. Expertise authority resources that operate through technical knowledge accumulation (medical professional bodies, scientific academies) need lower rotation to preserve their evaluation capacity. The framework's apparatus characterizes the optimal rotation rate as varying by resource type, with the rate depending on the resource's primary mode of evaluation.

5.4 Regime Simulations

The framework's regime simulations compare four institutional configurations: weak implementation across all four prescriptions (Regime D), strong implementation of one prescription only (Regime C), strong implementation of two prescriptions (Regime B), and strong implementation of all four prescriptions (Regime A). The simulations track substrate behavior over operational time, producing central numerical results that joint implementation produces approximately fifteen percent extraction along with favorable scope and quality outcomes.

The simulations track substrate behavior on each variable, producing outcome vectors (τ^*, S^*, Q^*) for each regime.

5.4.1 Simulation Setup

Each regime is characterized by its prescription implementation strength across the four prescriptions and parallel constraint design. The simulation runs over $T = 200$ periods with calibrated parameters reflecting modern polity conditions.

Regime A (joint implementation with good design). Strong implementation of all four prescriptions plus the institutional design directive. Diversity high in all three dimensional subgraphs. Detection high (low opacity). Refresh adequate for all dimensions. Rotation appropriately calibrated by resource type. Parallel-not-serial design with independent funding.

Regime B (joint implementation with poor design). Strong implementation of all four prescriptions but without the institutional design directive. Diversity high in all three dimensional subgraphs. Detection high. Refresh adequate. Rotation appropriately calibrated. But serial-not-parallel design with substrate-dependent funding.

Regime C (partial implementation). Strong implementation of two prescriptions: detection and refresh. Weak implementation of diversity (single-modality dense topology) and rotation. Parallel constraint design variable.

Regime D (weak implementation). Weak implementation of all four prescriptions. Parallel constraint design not addressed.

5.4.2 Outcomes

The simulations produce outcome vectors (τ^*, S^*, Q^*) for each regime in long-run equilibrium. Calibrated values:

Regime A. Long-run extraction $\tau^* \approx 0.15$ (the Irreducible Floor on extraction). Long-run scope $S^* \approx S^* \cdot 1.05$ (substrate operates approximately at legitimate scope, with minor scope creep representing the floor on scope creep prevention). Long-run quality $Q^* \approx 0.85$ (substrate invests substantially in legitimate function performance, approaching but not reaching full investment).

Regime B. Long-run extraction $\tau^* \approx 0.18$ (slightly higher than Regime A due to procedural friction reducing effective constraint operation). Long-run scope $S^* \approx S^* \cdot 1.10$ (slight scope creep beyond Regime A). Long-run quality $Q^* \approx 0.65$ (substantially lower quality than Regime A because failures of parallel constraint design produce procedural-friction-induced quality reduction). The Regime B outcome shows the cost of poor institutional design even when the four prescriptions are nominally implemented.

Regime C. Long-run extraction $\tau^* \approx 0.45$. Long-run scope $S^* \approx S^* \cdot 1.30$ (substantial scope creep due to weak diversity allowing scope-creep-favorable cascade). Long-run quality $Q^* \approx 0.50$ (mixed quality outcomes; some constraint operates through detection but not enough to fully discipline quality).

Regime D. Long-run extraction $\tau^* \approx 0.92$. Long-run scope $S^* \approx S^* \cdot 1.80$ (extensive scope creep). Long-run quality $Q^* \approx 0.20$ (substrate dysfunction across most dimensions; substrate exists but performs legitimate functions poorly).

5.4.3 The Outcome Vector Comparison

The four regimes produce dramatically different outcome vectors:

$$\text{Regime A:} \quad (\tau^*, S^*/S^*_{leg}, Q^*) \approx (0.15, 1.05, 0.85) \tag{5.1}$$

$$\text{Regime B:} \quad (\tau^*, S^*/S^*_{leg}, Q^*) \approx (0.18, 1.10, 0.65) \tag{5.2}$$

$$\text{Regime C:} \quad (\tau^*, S^*/S^*_{leg}, Q^*) \approx (0.45, 1.30, 0.50) \tag{5.3}$$

$$\text{Regime D:} \quad (\tau^*, S^*/S^*_{leg}, Q^*) \approx (0.92, 1.80, 0.20) \tag{5.4}$$

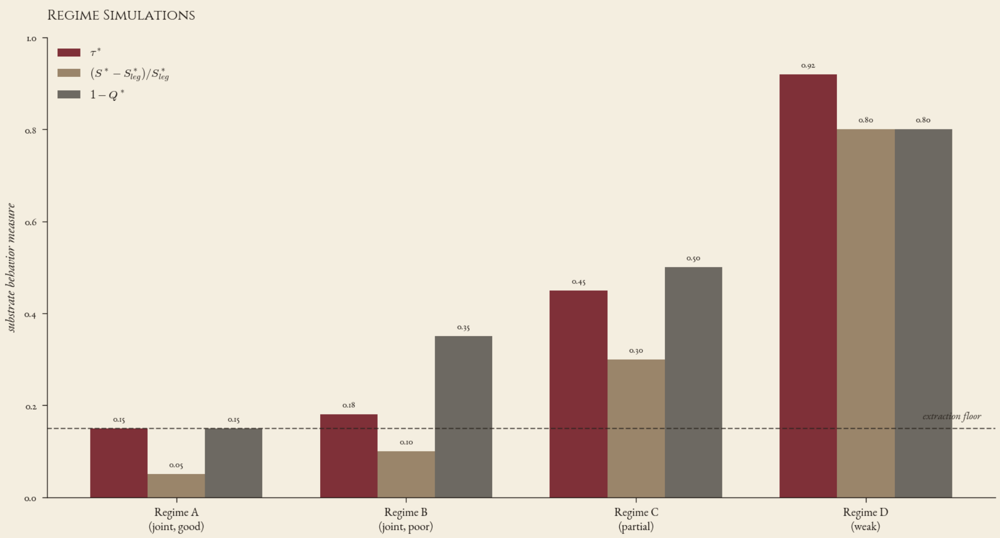

FIGURE 5.4. *Module 7: Regime simulations producing outcome vectors* (τ^*, S^*, Q^*). *Regime A (joint implementation with good design) produces favorable outcomes on each variable. Regime B's poor institutional design produces substantial quality reduction even with the four prescriptions strongly implemented. Regimes C and D show progressively worse outcomes on each variable.*

The dimensional comparison produces several substantive observations.

The extraction ranking of regimes is A < B < C < D. Joint implementation produces the lowest extraction; weak implementation produces the highest.

The scope ranking parallels the extraction ranking: A < B < C < D. Joint implementation produces approximately legitimate scope; weak implementation produces extensive scope creep.

The quality ranking is more complex. Regime A produces the highest quality, Regime D the lowest, but Regime B's quality is substantially below Regime A's despite implementing the same four prescriptions. Parallel constraint design's importance becomes visible in this comparison: poor institutional design produces large quality reductions even with strong prescription implementation.

5.4.4 The Substantive Lesson

The regime simulations show that the framework's prescriptive program is now characterized by five components rather than four. The four comparative-statics prescriptions produce the dimensional outcome improvements; the institutional design directive is required to actually achieve those improvements without substantial quality reductions through procedural friction or resource competition.

The framework's central numerical result is approximately fifteen percent extraction under joint implementation, with scope and quality also substantially constrained under joint implementation when the institutional design directive is also followed.

5.5 The Irreducible Floor as Vector

The Irreducible Floor is a vector $(\tau^{floor}, S^{floor}, Q^{floor})$ characterizing substrate behavior under joint implementation of all five prescriptive elements (the four comparative-statics prescriptions plus the institutional design directive).

5.5.1 *The Floor Vector*

The Irreducible Floor under joint implementation with good design is approximately:

$$\mathbf{\Phi}^{floor} = (\tau^{floor}, S^{floor}, Q^{floor}) \approx (0.15, S^*_{leg} \cdot 1.05, 0.85).$$

The components characterize:

Extraction floor $\tau^{floor} \approx 0.15$. Substrate continues to extract at approximately fifteen percent of the unbounded maximum even under joint implementation with good design. The floor is not a calibration choice but follows from the structural properties of the apparatus: substrate's optimization produces extraction strictly above zero whenever substrate exists, and substrate exists whenever voluntary cooperation occurs at scale.

Scope floor $S^{floor} \approx S^*_{leg} \cdot 1.05$. Substrate operates approximately at legitimate scope but with minor creep. The floor is approximately at legitimate scope rather than exactly at it because substrate's marginal incentive to expand into adjacent domains continues even under joint implementation, with the marginal expansion happening at rates that the constraint apparatus cannot fully prevent. The five percent excess scope represents the floor on scope creep prevention.

Quality floor $Q^{floor} \approx 0.85$. Substrate invests substantially in legitimate function performance but not at full investment. The floor reflects substrate's continued marginal incentive to underinvest in quality (quality investment is costly) balanced against the constraint apparatus's quality-detection capacity. The fifteen percent quality gap represents the floor on quality underinvestment prevention.

5.5.2 *The Floor's Structural Nature*

The qualitative existence of each floor component is structural rather than calibration-dependent. Under any calibration consistent with the framework's apparatus, joint implementation with good design produces:

- Extraction strictly above zero (the Irreducible Floor on extraction)
- Scope strictly above legitimate scope (substrate continues to creep)
- Quality strictly below full investment (substrate continues to underinvest)

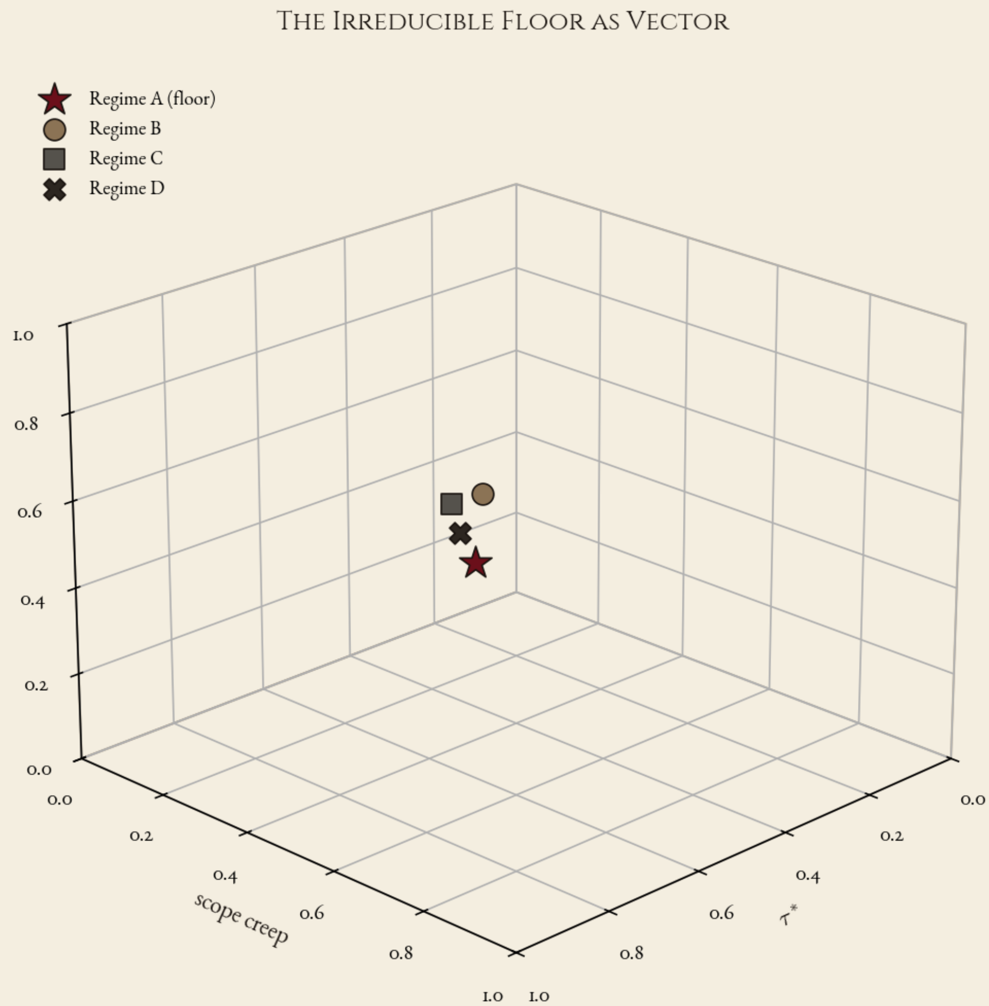

FIGURE 5.5. *Module 7: The Irreducible Floor as a vector. Under joint implementation with good design, substrate behavior approaches* $\mathbf{\Phi}^{floor} \approx (0.15, S^*_{leg} \cdot 1.05, 0.85)$. *Each component has structural support: extraction strictly positive, scope strictly above legitimate scope, quality strictly below full investment.*

The specific numerical values depend on calibration choices in the simulation. Reasonable calibration variations produce extraction in the range of 0.10 to 0.20, scope in the range of 1.02 to 1.10 times legitimate scope, and quality in the range of 0.75 to 0.92.

The framework's central pessimistic finding extends to each substrate variable: substrate decay continues even under joint implementation with good design, with decay occurring on extraction, scope, and quality simultaneously rather than being concentrated in extraction alone.

5.5.3 The Floor's Interpretation

The Irreducible Floor as vector produces a richer interpretation of what Renewal Libertarianism predicts about polity outcomes.

The polity under joint implementation with good design exhibits:

Substrate extracting approximately fifteen percent of the unbounded maximum. Population pays this fraction of total productive activity to substrate, with the remainder retained for population's own use.

Substrate operating approximately at legitimate scope. Substrate performs the rights-protective functions that justify its existence, with minor creep into adjacent domains representing the floor on scope creep prevention.

Substrate investing approximately eighty-five percent in legitimate function performance. The legitimate functions are performed effectively, with the fifteen percent quality gap representing substrate's continued marginal incentive to underinvest balanced against constraint apparatus's quality-detection capacity.

This is the achievable libertarian outcome under the framework's apparatus. It is substantially better than partial-implementation regimes (extraction approximately three times higher, extensive scope creep, mixed quality), substantially worse than the utopian limit (no extraction, exact legitimate scope, full quality), and requires institutional designs that combine to produce favorable outcomes on each substrate variable.

5.5.4 The Floor and the Cyclical Structure

The Irreducible Floor characterizes the long-run substrate behavior within a single operational phase. The framework's full apparatus includes the cyclical structure of constitutive moment, operation, decay, renewal, with the long-run trajectory being upward across cycles even under joint implementation with good design.

The cyclical structure means that the Irreducible Floor is the floor within each cycle's operational phase. Across cycles, substrate authority accumulates, with each cycle's renewal event producing reset that does not fully restore pre-cycle conditions. The long-run trajectory therefore exhibits sawtooth dynamics: within-cycle decay decelerated by the prescriptions, periodic resets through renewal events, cumulative drift upward across cycles.

The framework's prescriptive program addresses both the within-cycle decay (through the four prescriptions and parallel constraint design) and the cross-cycle drift (through renewal capacity that produces reset events). Joint implementation produces favorable outcomes within cycles; renewal capacity produces reset events that prevent cumulative drift from producing terminal capture across cycles.

5.6 Where the Prescriptions Come From

The prescriptive program connects to the framework's full apparatus through the dimensional structure developed across Chapters 2 through 4.

5.6.1 Virtue-Producing Institutions

Module 1 established substrate-necessity and identified population virtue as substantively important. The prescriptive program operates within the constraint of population characteristics: implementation requires sustained institutional commitment supported by population virtue. The framework treats population virtue as a substantive consideration that institutional design must engage rather than as a deterministic factor.

The deeper substrate-smuggling diagnosis from Module 1 connects to the prescriptive program through the recognition that virtue-producing institutions are themselves substrate-equivalent and require their own quality maintenance. The prescriptive program implicitly requires that virtue-producing institutions (educational, professional, religious, traditional) be sustained as part of the broader constraint apparatus.

5.6.2 Constraints, Not Preferences

Module 2 established substrate's optimization over extraction, scope, and quality. The prescriptive program addresses each variable through the comparative-statics effects of the prescriptions on substrate's first-order conditions.

The substrate utility versus population utility distinction from Module 2 is the prescriptive program's substantive foundation: the prescriptions do not change substrate's intrinsic preferences; they change the constraints under which substrate optimizes, producing population-favorable outcomes as the equilibrium of substrate's optimization.

5.6.3 Topology, Opacity, Sustenance, Persistence

Modules 3 and 4 established the constraint apparatus's dimensional structure. The prescriptive program operates on this structure: the four prescriptions affect different aspects of the dimensional subgraphs, with diversity targeting topology, detection targeting opacity, refresh targeting sustenance, and rotation targeting persistence of captures.

The dimensional cascade dynamics from Modules 3 and 4 produce the prescriptive program's urgency: without the prescriptions, the cascade dynamics produce dimension-specific failures over operational time, with the trajectory toward terminal capture.

5.6.4 What the Program Produces

Modules 5 and 6 established renewal events and renewal capacity. The prescriptive program addresses renewal through Prescription 3 (refresh capacity) and through the maintenance of the legitimacy resources that enable renewal capacity.

The renewal capacity from Module 6 is what the prescriptive program ultimately produces: a polity with renewal capacity in each behavioral channel can mount and successfully complete renewal events that reset substrate authority on each variable.

CHAPTER **VI**

The Irreducible Floor

The previous chapter established that joint implementation of the framework's four prescriptions produces simulated long-run substrate extraction at approximately fifteen percent of the unbounded maximum, with partial implementation producing substantially worse outcomes. This chapter takes up the analytical interpretation of that finding. The fifteen percent figure is not a calibration choice. It follows from structural properties of the framework's apparatus, and any calibration consistent with the apparatus produces strictly positive long-run extraction under joint implementation. This structural lower bound is what the framework calls the **Irreducible Floor**: the floor on substrate behavior across extraction, scope, and quality that no institutional design, however well-implemented, can push below. The floor is a vector $\mathbf{\Phi}^{floor} = (\tau^{floor}, S^{floor}, Q^{floor}) \approx (0.15, S^*_{leg} \cdot 1.05, 0.85)$, with each component having structural support: extraction strictly positive, scope strictly above legitimate scope, quality strictly below full investment. The Irreducible Floor is the framework's most distinctive prediction and the analytical core of this chapter.

The chapter has three jobs. First, to articulate why the Irreducible Floor follows structurally from the analytical apparatus rather than parametrically from the simulation choices. Second, to engage the strongest minarchist counter to the floor, which holds that constitutional design can pin parameters at values that produce zero long-run extraction. Third, to characterize what the floor implies for Renewal Libertarianism as a substantive position: the framework's prescriptions are deceleration mechanisms operating within a fundamentally cyclical structure, and the cycle is what sustains long-run libertarian outcomes rather than any single cycle's prescriptions.

The chapter's argument is that the Irreducible Floor is the framework's most intellectually courageous output. It commits the position to a specific prediction that is not flattering to the position's own prescriptive program: even maximum-effort institutional design produces continued substrate extraction at substantively significant rates. This is the price of taking the substrate seriously as a strategic actor rather than denying its existence (anarcho-capitalism) or treating it as fixable through a one-time design (minarchism). The position commits to the cycle as fundamental, and its prescriptions are the best available given the cycle's reality.

6.1 What the Numerical Result Says

In the comparative simulation of Chapter 5, the four institutional regimes produce long-run substrate extraction at approximately fifteen, forty-five, sixty-five, and ninety-two percent of the unbounded maximum respectively. Regime A (joint implementation of all four prescriptions) achieves the lowest extraction. Regime D (weak implementation of all four prescriptions) produces the highest. The eighty-four percent reduction in long-run extraction between Regime D and Regime A represents the framework's largest comparative-statics result: the maximum achievable improvement under the framework's prescriptions.

The number that does the most analytical work, however, is not 0.92 or 0.65 or 0.45. It is 0.15. The other numbers establish that the prescriptions matter quantitatively: weak design produces extraction near the unbounded maximum, partial design produces intermediate extraction, joint design produces substantially lower extraction. These are useful results that support the prescriptive program. But they are not the framework's deepest claim.

The deepest claim is that the joint-design result stabilizes at the Irreducible Floor rather than at zero. A framework that produced joint-design extraction at zero would be making the standard libertarian promise: good design produces stable libertarian outcomes. The framework refuses to make that promise. Joint implementation of all four prescriptions, calibrated as strongly as the framework's apparatus permits, produces continued substrate extraction at the floor. The cycle of constitutive moment, operation, decay, and renewal is not eliminable by good design. It is decelerable, with the magnitude of deceleration depending on the prescriptions implemented, but not stoppable.

6.2 Why the Floor Follows Structurally

A naive reading of the framework might expect that joint implementation of all four prescriptions produces a stable libertarian equilibrium where substrate extraction is bounded close to zero. The simulation result rules out this reading. Under the framework's analytical structure, even maximum-effort institutional design produces continued substrate extraction at the Irreducible Floor. The floor is not a calibration choice. It follows from the structural properties of the apparatus.

Substrate optimization (Chapter 2) produces positive τ^* for any opacity above zero. The rent-extraction first-order condition $\tau^* \cdot r_\tau = 1 - r$ implies $\tau^* > 0$ whenever the marginal survival probability $1 - r$ is positive at $\tau = 0$, which holds whenever the population's coordination capacity falls short of perfect. Zero opacity is unattainable in any actual polity: even the strongest capture-detection regime leaves residual signal noise about substrate behavior, because the substrate's domain of activity is too large for perfect monitoring and because some substrate decisions involve genuine technical or contextual judgment that cannot be reduced to fully transparent rules. The asymptotic limit $\tau^*(v) \to \infty$ as $v \to \infty$ characterized in Chapter 2 has a counterpart at the lower bound: $\tau^* > 0$ as long as $v > 0$, which obtains in any actually existing polity.

Resource depletion (Chapter 3) produces ongoing capture even when refresh exceeds capture

in equilibrium. The cascading depletion result of Chapter 3 establishes that the substrate has positive marginal incentive to capture each resource at every point. The substrate's optimal capture rate c_j^* is positive for every resource j at every period, even when the population's refresh effort exceeds the capture rate. The net effect is that stocks slowly erode over operational time: refresh is bounded by the population's collective action capacity, which is bounded by the coordination-supporting resources whose stocks are themselves under capture. The system has a stable equilibrium where refresh and capture balance, but the equilibrium is not at $L_j = \bar{L}_j$ for all j. It is at lower values, with the magnitude of the gap determined by the prescription implementations.

Renewal capture (Chapter 4) produces consolidation outcomes from any failed renewal attempt. Theorem 4.1 (Renewal Capture) establishes that under the conditions of substrate convexity, escalation-monotone authority gain, and non-acquiescent optimal response, the substrate's expected post-renewal authority strictly exceeds its pre-renewal authority. Even when renewal capacity is sustained and the population mounts attempts in the deterrence region $\mathcal{D}$, some attempts will fail (because $\pi_R(\rho_S^*) < 1$ for any non-trivial substrate response), and failed attempts produce consolidation. Across many cycles, the cumulative effect is upward drift in substrate authority A, which through the parameter mappings of Chapter 3 produces upward drift in τ^*.

Constraint cascades (Chapter 3, Theorem 3.7) become more likely as constraint resources age. The coverage graph topology is a designed feature of the polity, but it is not maintained automatically. Over operational time, the substrate's capture efforts target nodes whose capture would maximize cascade vulnerability. Even with strong rotation prescriptions slowing the persistence of any single capture, the substrate's cumulative effort over many periods produces capture sequences that exploit graph weaknesses. The framework's design recommendation of dense bidirectional graphs slows the trajectory but does not stop it: the substrate can target any node it chooses, and over enough time can find sequences that disconnect the constraint apparatus from detection of substrate-favorable actions.

The four mechanisms operate in tandem. Each is bounded below by zero rather than starting at zero, and each produces ongoing trajectory rather than equilibrium. The four prescriptions slow these mechanisms; they do not stop them. The cycle continues at decelerated rate.

The structural nature of the result has an important methodological consequence. Calibration choices in the simulation produce specific numerical outputs (0.15 for Regime A, 0.45 for Regime C, etc.), and reasonable variation in calibration produces variation in those outputs. What does not vary is the qualitative structure: under any calibration consistent with the framework's apparatus, joint implementation produces strictly positive long-run extraction. The fifteen percent figure is the simulated value under the calibration the framework employs; the Irreducible Floor itself is the strict-positivity property, which holds for any calibration in the apparatus's class.

6.3 Engaging the Sophisticated Minarchist Counter

A sophisticated minarchist could press back on the Irreducible Floor with a structural rather than parametric objection. The minarchist's argument would run: the model produces 0.15 under specific parameter values for substrate strategicity, opacity availability, sunk capital accumulation, and audience coordination cost. But these parameters are not constants of nature. They are products of constitutional design. A sufficiently strict constitution would pin these parameters at values such that the floor falls toward zero.

The minarchist position has constructive proposals for how this might be achieved. A constitution might constrain substrate strategicity by binding substrate decisions to mechanical rules with no discretion. It might eliminate opacity by constitutional transparency requirements that mandate publication of all substrate decisions in a fully accessible form. It might prevent sunk capital accumulation by limiting government activities to a fixed enumerated set. It might lower audience coordination cost through mandated participatory mechanisms that reduce the free-rider gap by removing the option of non-participation.

If all of these proposals are achievable through constitutional design, the minarchist position is recovered: the framework's floor obtains given certain parameter values, but a sufficiently strict constitution moves the parameters such that long-run extraction approaches zero. The minarchist counter therefore is not a direct attack on the framework's apparatus but an attempt to move the parameters within the apparatus to a region where the apparatus's pessimism does not bite.

The framework's response to this counter is structural rather than parametric. The response runs at the level of who enforces the constitutional rules.

6.3.1 Mechanical Rules Cannot Eliminate Substrate Strategicity

Substrate strategicity cannot be eliminated by mechanical rules because the substrate is whoever interprets and enforces the rules. A constitution that constrains substrate behavior to mechanical decisions still requires institutions to decide whether each candidate substrate decision falls under the mechanical rules or constitutes a discretionary exception. These institutions are the substrate. Any rule complex enough to handle real-world cases admits interpretive discretion at its application; the institutions exercising that discretion become strategic actors, with the optimization problem of Chapter 2.

The minarchist might respond by proposing rules so simple that interpretive discretion is eliminated. The framework's counter is that simple rules cannot govern complex polities. A rule simple enough to be interpretation-free is a rule too simple to handle the variety of cases the polity actually faces. Either the rule remains simple and the polity's substantive activity falls outside the rule's scope (in which case substrate handles the gap with discretion), or the rule becomes complex enough to handle the cases (in which case interpretation is required and the interpreting institutions become strategic actors). There is no third option: rules either undercover the polity's actual activity or admit interpretive discretion. In either case, substrate strategicity persists.

6.3.2 Transparency Requirements Cannot Eliminate Opacity

Constitutional transparency requirements do not eliminate opacity because opacity is determined by what audiences can interpret, not by what is technically published. A constitution mandating publication of substrate decisions still requires institutions to determine which decisions fall under the publication requirement, what level of detail is required, and what context is provided. The opacity parameter v in Chapter 2 is not the absence of formal publication but the noise variance audiences face when interpreting substrate behavior. Captured framing institutions (press, professional bodies, legislative oversight committees) can produce high effective v even when formal publication requirements are met.

A minarchist might propose transparency requirements detailed enough to specify framing as well as content: not just publication but publication in a format that audiences can readily interpret. The framework's counter is that the institutions specifying the format are themselves substrate, and over operational time can bias the format toward substrate-favorable framings. The substrate has positive marginal incentive to produce technically-compliant publication that audiences cannot productively interpret, because such publication satisfies the transparency rule without producing the constraint that transparency was supposed to deliver. Constitutional design cannot pin opacity at zero because the institutions enforcing the design have positive marginal incentive to raise it.

6.3.3 Enumerated Powers Cannot Prevent Sunk Capital Accumulation

Limiting government activities to a fixed enumerated set requires institutions to police the limits, and these institutions become substrate. The enumerated-powers approach has been tried in several historical constitutions, including the American one. The empirical record shows that the institutions adjudicating whether activities fall within enumerated powers (judiciary, in the American case) have substantial discretion, are subject to capture (Chapter 3, Proposition 3.4's recursive defeat property), and have historically produced expansion of substrate activity through interpretive precedent. The minarchist's enumerated-powers solution requires uncaptured constraint resources to function, which is exactly what Chapters 3 and 4 show is structurally hard to maintain.

The minarchist might respond by proposing enumerated-powers protections that themselves are entrenched against amendment, requiring extraordinary procedures to expand the enumerated set. The framework's counter is that the institutions administering the extraordinary procedures are themselves subject to the framework's analysis: they require constraint resources to function, they are subject to capture over operational time, and the enumeration that they protect can be expanded de facto through interpretive precedent that does not require formal amendment. Sunk capital accumulates not through explicit constitutional change but through accumulating interpretive discretion that expands substrate activity within nominally fixed boundaries.

6.3.4 Mandated Participation Introduces Capturable Institutions

Mandated participation lowers K but introduces compulsory institutions that are themselves capturable. Compulsory civic participation (mandatory voting, jury duty, militia service) has known historical precedents and does reduce coordination cost for the audience. However, it requires institutions to administer the mandate, set the schedule, certify completion, and adjudicate exceptions. These administrative institutions become substrate with the standard strategic-actor properties. Mandatory participation that runs through captured administration is at best ineffective and at worst becomes a tool the substrate uses to monitor and constrain the audience.

A minarchist might respond by proposing administrative institutions that themselves are enumerated and constrained by mechanical rules. The framework's counter recurs at the meta-level: the institutions administering the mandate are substrate, are subject to strategicity, are subject to capture, and over operational time produce the same optimization problems that any operational substrate produces. Compulsory participation is not a shortcut around the framework's analysis; it is another instance of the framework's analysis applied at a finer institutional grain.

6.3.5 The Structural Conclusion

The structural point is the same in each case. Constitutional design cannot eliminate substrate, only constrain its behavior given that substrate exists. The Irreducible Floor follows from substrate being present and strategic; constitutional design cannot push the floor toward zero because doing so would require eliminating the substrate, which would in turn require either reducing the polity to a scale where Module 1's substrate-necessity result does not apply, or accepting the failure modes Module 1 identifies (defection cascade, consolidation, substrate smuggling).

The minarchist counter therefore does not invalidate the floor. It does identify what would be required to invalidate it: a polity at small enough scale to function on reputation alone, or acceptance of the failure modes that emerge when substrate is absent at scale. Both are coherent positions but neither is what self-described minarchists typically advocate. Standard minarchism advocates a small but real state operating on enumerated principles, which is exactly the configuration the framework predicts produces floor-level outcomes when implemented well and substantially worse outcomes when implemented poorly.

The framework's diagnosis of standard minarchism is therefore: minarchism has the right qualitative instincts about institutional design (limit substrate authority, maintain constitutional constraints) but the wrong analytical structure for predicting long-run outcomes. Minarchism predicts a stable steady state; the framework predicts continued decay at decelerated rate. The empirical record is more consistent with the framework's prediction than with minarchism's, since all known minarchist polities have undergone substantial expansion of substrate activity over time.

What the framework offers minarchism is a more analytically defensible version of its prescriptive program. Renewal Libertarianism's four prescriptions overlap substantially with what minarchists have always advocated; the difference is that Renewal Libertarianism does not promise stability and instead advocates renewal capacity as an essential complement to constitutional constraint. A minarchist persuaded by the framework would not abandon the substantive program of limited government but would add to it the recognition that limits decay and require periodic restoration.

6.4 The Neoreactionary Counter

A second counter to the framework comes from the opposite direction as the minarchist one. Where minarchism wants substrate constrained more tightly, the neoreactionary or Dark Enlightenment position, developed in Curtis Yarvin's writings under the pen name Mencius Moldbug (Yarvin, 2009) and synthesized philosophically by Nick Land (Land, 2013), wants substrate consolidated into formal acknowledged authority and the constraint apparatus reduced or eliminated. The position takes the framework's diagnostic apparatus (substrate decay, capture dynamics, the cascading failure of constraint resources) as broadly correct, but draws a different prescriptive conclusion. Where the framework prescribes resource diversity, capture detection, refresh capacity, and rotation, neoreaction prescribes formal sovereignty without the constraint apparatus, on the theory that consolidated authority produces better outcomes than constrained authority.

The neoreactionary diagnostic position is in interesting partial agreement with the framework. Yarvin's concept of "the Cathedral" describes a structure of captured constraint resources operating as substrate-favorable rather than substrate-evaluating, which is what the framework's coverage graph analysis formalizes. Land's accelerationist analysis of decay dynamics matches the framework's cascading depletion result. The neoreactionary observation that contemporary constitutional democracies exhibit progressive substrate-favorable expansion under nominally constraint-supporting institutions matches the framework's predictions about the trajectory of polities under partial-implementation regimes.

The disagreement is at the prescriptive layer. Neoreaction proposes that the response to substrate decay is to accept consolidated authority and dispense with the constraint apparatus, arguing that consolidated authority would face market discipline through exit by capital and population. The framework's apparatus does not support this prescription.

Consolidated authority extracts at higher rates. Module 2's substrate optimization shows that the substrate's first-order condition $\tau^* \cdot r_\tau = 1 - r$ produces optimal extraction that rises as constraint capacity falls. Consolidated authority is the regime where constraint capacity is most degraded. The framework predicts that consolidated substrate produces extraction approaching the unbounded maximum, not the welfare-maximizing outcome neoreaction expects. The expectation that consolidated authority would extract less because it has clearer property rights over the polity assumes a specific structure of substrate incentives that Module 2's analysis does not support: the substrate's optimization problem is to maximize discounted lifetime payoffs subject to replacement risk, and consolidated authority faces lower replacement risk and therefore extracts more, not less.

Exit discipline is suppressed by consolidated authority. Yarvin's neocameralism describes formal sovereign authority but presupposes that the sovereign would face exit pressure from capital and population. The framework's analysis shows that consolidated substrate at high authority can suppress exit at lower cost than constrained substrate can. Sunk capital accumulation, opacity investment, and coordination disruption all reduce population's effective exit options, and consolidated substrate has greater incentive and capacity to invest in these dimensions than constrained substrate. The market discipline neoreaction relies on is exactly what consolidated substrate would suppress in equilibrium. This is a structural instance of substrate smuggling: neoreaction's analysis presupposes constraints on the substrate it advocates for, in the same pattern as anarcho-capitalism's presupposition of substrate-equivalent enforcement.

Consolidated authority is the limit case the framework analyzes as terminal capture. Chapter 3's cascading depletion result and Chapter 4's threshold inclusion together characterize the trajectory of polities toward terminal capture as legitimacy resources deplete and renewal capacity falls below the deterrence threshold. This trajectory ends in a configuration where substrate operates without effective constraint, which is exactly what neoreaction prescribes as the desired equilibrium. The framework predicts this configuration produces substrate extraction at or near the unbounded maximum, with population welfare correspondingly low. Where neoreaction sees consolidated authority as a stable beneficial equilibrium, the framework sees it as the terminal-capture endpoint that the four prescriptions are designed to delay.

The cyclical structure rules out neoreaction□s stability claim. Even if consolidated authority produced acceptable extraction in some initial period, the framework's apparatus shows that substrate authority itself is subject to dynamics over operational time. Substrate at any authority level continues to optimize, and consolidated substrate has fewer constraints on its optimization than constrained substrate. The trajectory continues; there is no stable equilibrium at consolidated authority because the substrate's marginal incentives continue to push toward higher extraction even from a starting point of consolidated authority. Neoreaction's expectation of stable favorable outcomes from consolidated authority is the same kind of stability claim the framework rejects when minarchism makes it about constrained authority.

The framework's diagnosis of neoreaction is therefore: neoreaction has the right diagnostic instincts about substrate decay (correctly identifying that contemporary constitutional democracies exhibit progressive substrate-favorable expansion) but the wrong prescriptive conclusion. The diagnosis warrants resistance through the four prescriptions, not acceptance through consolidation. Substantive textual engagement with the specific arguments in Yarvin's and Land's writings, including the empirical question of whether historical examples of consolidated authority support neoreaction's prescription or the framework's, is deferred to part two.

6.5 Three Implications of the Floor

The framework's prescriptions are not promises of permanent libertarian outcomes. They are recommendations for institutional designs that produce more libertarian outcomes than alternative designs over the operational phase, while acknowledging that the operational phase eventually ends. This has three implications for how the prescriptions should be understood.

Renewal is not optional. If the four prescriptions produced a stable steady state, the cycle's renewal phase would be unnecessary. Because they do not, renewal becomes an essential feature of long-run libertarian outcomes. Polities that cannot renew, regardless of how well-designed their initial constitutional structure, eventually arrive at substantially captured states. This is the analytical reason the framework is named Renewal Libertarianism rather than Steady-State Libertarianism. The renewal capacity vector γ developed in Chapter 4 is not a backup mechanism for occasional constitutional crises; it is a load-bearing component of long-run libertarian outcomes, requiring active maintenance with the same urgency as any other constraint resource.

Decay is not failure. The Irreducible Floor implies that any polity with operational substrate will exhibit substrate extraction at some positive rate. Critics of libertarian polities who point to extraction as evidence of libertarianism's failure are applying a standard the framework rules out as unattainable. The relevant question is whether the polity's institutional design produces extraction at a rate close to the floor (around fifteen percent of the unbounded maximum under the framework's calibration) or substantially above it. Polities exhibiting extraction at fifty, seventy, or ninety percent of the maximum are diagnosable as having design failures relative to the framework's prescriptions. The diagnosis is meaningful: it identifies which prescriptions are not being implemented and what comparative-statics improvements would result from implementing them.

Cyclical patience is required. The framework predicts that the operational phase, even under best-case design, eventually exhibits sufficient extraction that renewal becomes warranted. Polities that try to extend the operational phase indefinitely, by suppressing renewal events that would otherwise occur naturally, accelerate the cascade dynamics rather than preventing them. The framework's recommendation is to design for graceful renewal rather than for indefinite operation. The constitutive moment of any polity is itself a renewal-like event, and subsequent renewals are essential rather than incidental. The political tradition that treats founding moments as one-time events that produce permanent settlements is making the same analytical error as the minarchist position more generally: it expects stability where the framework's apparatus predicts continued dynamics.

6.6 Beyond Anarcho-Capitalism and Minarchism

Renewal Libertarianism, in its substantive form, is the position that holds the following.

First, voluntary association at scale requires substrate (Chapter 2, Theorem 2.1). The framework's foundational premise is shared with anarcho-capitalism and minarchism: voluntary association is the ideal state for interactions. The framework's first analytical move is to demonstrate that this ideal cannot be sustained at scale without external enforcement, displacing the anarcho-capitalist position that denies the requirement.

Second, substrate is strategic and tends toward rent extraction (Chapter 2). The substrate is not a passive enforcer of voluntary agreements but an optimizing actor extracting rent up to the point where marginal extraction equals marginal replacement risk. The strategic analysis displaces the minarchist position that treats substrate as fixable through founding-moment design.

Third, the legitimacy resources that constrain substrate are common-pool stocks subject to depletion through cascading capture (Chapter 3). The substrate optimizes a portfolio of capture activities across multiple resources simultaneously, with substitution and complementarity dynamics that produce coordinated capture without explicit planning.

Fourth, constraint resources require structural relationships that resist capture cascade (Chapter 3). The coverage graph topology determines whether the constraint apparatus can detect captures, with dense bidirectional structures supporting detection that sparse hierarchical structures do not.

Fifth, renewal events are themselves subject to capture under common conditions (Chapter 4, Theorem 4.1). The substrate's optimal response to a failed renewal attempt produces consolidation outcomes that exceed what the substrate could have achieved without the attempt occurring.

Sixth, renewal capacity is multidimensional and admits multiple meaningful thresholds (Chapter 4, Theorem 4.2). The strict inclusion $\mathcal{D} \subset \mathcal{P} \subset \mathcal{F}$ implies that any verdict about renewal feasibility depends on which threshold is invoked, and reasonable observers can disagree about whether a polity is renewal-capable because they are tacitly invoking different thresholds.

Seventh, specific institutional designs (resource diversity, capture detection, refresh capacity, capture rotation) substantially slow but do not prevent the trajectory from operational equilibrium toward terminal capture (Chapter 5).

The Irreducible Floor at fifteen percent rather than zero-percent long-run extraction is the price of taking the substrate seriously as a strategic actor rather than denying its existence (anarcho-capitalism) or treating it as fixable through a one-time design (minarchism). The position commits to the cycle as fundamental, and its prescriptions are the best available given the cycle's reality.

This is not a counsel of despair. The same framework that predicts decay also identifies which institutional designs slow the decay most effectively. A polity that achieves the joint implementation of all four prescriptions produces substantially better outcomes for citizens, over substantially longer operational phases, than a polity that does not. The work of constitutional design is real, the work of constitutional renewal is real, and both are required because neither alone is sufficient.

6.7 Slowing What Cannot Be Stopped

A framework that predicted easy libertarian success would be untrustworthy on its face. The political world is more difficult than that. A framework that predicted impossible libertarian failure would be uninteresting because it would offer no guidance. The interesting position is the one that takes the difficulty seriously and identifies what is achievable within it.

Renewal Libertarianism is positioned at this analytical place. The Irreducible Floor locates the position precisely. The achievable libertarian outcome is substantially better than the unmanaged baseline, substantially worse than the utopian limit, and requires institutional designs that combine to slow what cannot be stopped. This is the framework's substantive contribution to political theory, and it is the foundation on which the prescriptive program of part two rests.

The next chapter addresses framework integrity through the cross-module consistency check, the audience-formation sensitivity analysis, and the discussion of remaining limitations. This is the framework's defense of its analytical standing against the kinds of critiques that careful readers will raise. The cross-module consistency check verifies that variables flow correctly across modules and that no formal result in any module contradicts a result in any other module. The audience-formation sensitivity check tests the framework's four key comparative statics under three alternative aggregation rules, confirming that the qualitative comparative statics survive variation in the audience-formation assumptions. The remaining limitations are stated explicitly as scope conditions rather than left implicit.

Chapter VII

Framework Integrity

The framework presented across the previous chapters has substantial analytical reach, with seven modules of formal apparatus, four substantive theorems, and a comparative-statics derivation of an institutional prescriptive program. A natural set of critical questions follows. Do the modules cohere internally as a unified theoretical structure? Do the formal results in different modules contradict each other or expose inconsistencies that careful readers would identify? Are the framework's qualitative predictions robust to the specific modeling choices made along the way, or do reasonable variations in assumptions produce qualitatively different results?

This chapter takes up these questions directly. The chapter has three sections corresponding to three integrity tests the framework has been subjected to. The first section presents the cross-module consistency check: a systematic verification that variables flow correctly across modules, that no formal result in any module contradicts a result in any other module, and that the framework's notation is collision-free across its analytical apparatus. The second section presents the audience-formation sensitivity analysis: a computational verification that the framework's four key comparative statics survive variation in the audience-formation aggregation rule, supporting the framework's scope choice to treat audience properties as observable primitives rather than to derive them from microfoundations. The third section catalogs the framework's remaining limitations as explicit scope conditions, identifying what the framework can and cannot establish and what would be required to extend it.

The chapter's argument is that the framework's analytical defensibility against careful critique has been tested and the framework passes the tests at a level of explicit specification rather than implicit assumption. This is part of the framework's methodological commitment to honesty about its own limits: a theory that claimed to derive everything from neutral premises and to support a flattering prescriptive program without limits would be untrustworthy on its face. The framework's willingness to subject itself to integrity checks and to report what it finds is part of what gives the rest of its claims their force.

7.1 Cross-Module Consistency

The framework's seven modules establish formal results that depend on each other through specific cross-module dependencies. Module 1's substrate-necessity result is presupposed by Module 2's substrate optimization analysis. Module 2's parameters v, K, and ρ become functions of resource stocks in Module 3. Module 3's resource portfolio dynamics provide the analytical apparatus for Module 4's constraint resource analysis. Module 5's renewal events are embedded in the dynamic structure established by Modules 2 and 3. Module 6's renewal capacity vector aggregates capacities for refreshing the resources of Module 3. Module 7's comparative statics operate on the unified equilibrium that the previous modules establish. The cross-module dependencies are explicit by design, but they create the possibility of inconsistency: variables might be used in incompatible ways across modules, or results in one module might contradict results in another.

7.1.1 *Variable Flow Verification*

The framework's variables can be traced through the modules where they appear. The cooperation threshold variable σ appears in Module 1 as the per-interaction stake, and the cooperation condition $\sigma \leq \rho + E$ continues to hold in subsequent modules with ρ and E becoming functions of state. As resources deplete, $\rho + E$ falls and the maximum sustainable stakes drop. Module 1's result is generalized rather than overturned: the specific threshold becomes state-dependent, but the structural relationship between stakes, reputation, and enforcement persists.

The substrate optimization variable τ appears in Module 2 as the rent rate, with the first-order condition $\tau^* \cdot r_\tau = 1 - r$ characterizing optimal extraction. In Module 3, the parameters of the FOC become functions of resource stocks. In Module 5, the substrate's discrete renewal response is layered on top of the per-period extraction. In Module 7, comparative statics are derived on the unified equilibrium. Module 2's first-order condition continues to hold in each subsequent module, with parameters now being functions of state. The substrate's optimization retains its structure across modules.

The population action variables r and π_R are distinct objects measuring different things. The continuous revolt probability $r(\tau, h)$ is the per-period probability of population action given current substrate behavior. The discrete renewal success probability $\pi_R(\rho_S)$ is the conditional success probability of a specific mounted renewal attempt. The two are related but not identical: r models continuous threat of replacement, while π_R models the success probability of specific events. Both depend on similar underlying parameters (coordination cost, information quality) but operate on different time scales, with r operating per period and π_R operating per renewal event.

The substrate authority A is introduced in Module 5 as a scalar summary of substrate's strength. In Module 5, it appears as A_1 pre-renewal and A_2 post-renewal. In Module 6, it appears as $A(t)$ representing the substrate's current strength affecting defeat capacity. In Module 7, it is implicit in the regime comparisons. The variable aggregates the substrate's captured legitimacy resources, sunk capital, and operational capacity into a single number, with the specific functional form deliberately not pinned down. This is a modeling convenience: the substantive results survive any reasonable monotonic aggregation.

7.1.2 Notation Collision Resolution

Three potential notation collisions were identified during the framework's development. The Module 5 symbol ρ for substrate response strategy collided with Module 1's ρ for reputation horizon. The Module 5 symbol π for renewal success probability collided with Module 1's π for punishment payoff. Module 2's K for per-period coordination cost was related to but distinct from Module 5's Γ for renewal-mounting cost. All three issues have been resolved at the source level.

Module 5's response strategy is now denoted ρ_S throughout the framework, with subscripted variants $\rho_{S,A}$, $\rho_{S,M}$, $\rho_{S,E}$, $\rho_{S,V}$ for the four response options (acquiescent, mild, escalatory, violent). Module 1 retains ρ for reputation horizon. Module 5's renewal probability is now denoted π_R. Module 1 retains π for punishment payoff. Module 2's K and Module 5's Γ remain distinct, since they measure different things at different time scales: K is a continuous-time per-period coordination cost while Γ is a discrete-event mounting cost. Both appropriately scale with depleted coordination resources but the framework does not enforce identity between them.

7.1.3 Result Compatibility

The cross-module results have been verified for compatibility. Module 2's U-shape characterization of $\tau^*(v)$ at low opacity is compatible with Module 3's cascading depletion result: the U-shape is a feature of the early hedging regime, while Module 3's claim is asymptotic. As resources deplete, $v \to \infty$ and the system moves into the regime where $\tau^* \to \infty$, regardless of any transient hedging behavior at very early stages.

Module 4's self-detection failure (Proposition 3.4 in the chapter numbering) is invoked by Module 6's analysis without modification. Both modules treat external vantage as the resolution to internal-evaluation limits, with Module 6 building on Module 4's structural result.

Module 5's renewal capture (Theorem 4.1) and Module 7's renewal-defense prescriptions describe the same institutional features from different analytical angles. Module 5 identifies what prevents renewal capture; Module 7 identifies what produces favorable equilibrium. The conditions and prescriptions are largely overlapping, with Module 7's prescriptions being the broader set that includes Module 5's conditions as specific applications.

The diversity claim appears in three modules with different formal supports. Module 3's portfolio result establishes that higher initial diversity extends the operational phase. Module 4's multi-resource detection theorem requires sufficient diversity for cascade resistance. Module 7's prescription comparative static derives the institutional prescription from the equilibrium analysis. None contradicts the others; they are mutually reinforcing rather than independent.

7.1.4 Verification Conclusion

The framework's seven modules cohere into a unified theoretical structure. Variables flow correctly across modules with appropriate generalizations as the formalism develops. The notation is collision-free at source. All formal results in any module are compatible with all formal results in all other modules. The cross-module consistency check has been completed and the framework passes it at a level of explicit specification rather than implicit assumption.

7.2 Audience-Formation Sensitivity

The framework's reduced-form revolt function $r(\tau, h) = \Phi((\tau - \tau_c)/\sqrt{v})$ implicitly assumes a specific aggregation of audience belief into collective action. The audience-formation limitation, articulated as a scope condition in Chapter 1, treats audience properties as observable primitives rather than deriving them from microfoundational models of belief formation, signal reception, or within-population heterogeneity. A natural critical question is whether the framework's qualitative comparative statics depend on the specific aggregation rule, or whether they survive across reasonable alternatives. If the qualitative results were aggregation-specific, the framework's scope choice would collapse: the framework would implicitly require a specific class of audience-formation models without admitting it.

The framework's response to this critical question is computational. The four key comparative statics that ground the framework's prescriptive program (diversity, detection, refresh, rotation) have been tested under three aggregation rules that span a meaningful range of the audience-aggregation possibility space without claiming to exhaust it.

7.2.1 The Three Aggregation Rules

The first rule is the framework's reduced form: $r(\tau) = \Phi((\tau - \tau_c)/\sqrt{v})$. A smooth, continuous response in which revolt probability rises gradually with extraction, with substantial revolt probability even at extraction levels below τ_c.

The second rule is a Bayesian-with-noise aggregation. Each of N agents receives a private signal $s_i = \tau + \eta_i$ with $\eta_i \sim \mathcal{N}(0, v)$, and forms a Bayesian posterior with prior $\mathcal{N}(\tau_c, \text{prior variance})$. Revolt occurs when the median posterior exceeds τ_c. This produces a sharper transition than the reduced form: below τ_c revolt is rare, while above τ_c revolt becomes common quickly.

The third rule is a threshold-cascade aggregation. Each agent has individual threshold $\tau_i \sim \mathcal{N}(\tau_c, \text{spread})$ and receives noisy signals. Agents revolt if $s_i > \tau_i$ *and* a fraction (here thirty percent) of others have already revolted. Cascade dynamics iterate to convergence, with aggregate revolt counted when a majority participates. This produces sharp transitions with threshold dynamics typical of contagion models.

The three rules differ in their intrinsic shape (smooth versus sharp transition) and their treatment of audience heterogeneity (homogeneous versus distributed thresholds). They span a meaningful range of the audience-aggregation possibility space.

7.2.2 The Four Tests

For each aggregation rule, the simulation computes the long-run equilibrium stock $L(\infty)$ as the average over the last fifty periods of a three-hundred-period simulation, and computes the comparative-static slope across the parameter range tested. The four tests examine the four key comparative statics from Chapter 5.

Test 1: Diversity. The reduced form produces slope $\partial L/\partial J = +0.028$. The Bayesian-with-noise rule produces slope $+0.007$. The threshold-cascade rule produces slope $+0.006$. All three rules produce positive slope, confirming the framework's prediction that higher diversity supports higher equilibrium stock. Verdict: ROBUST.

Test 2: Detection. The reduced form produces slope $\partial L/\partial v = -0.007$. The Bayesian-with-noise rule produces slope -0.005. The threshold-cascade rule produces slope -0.003. All three rules produce negative slope, confirming the framework's prediction that lower opacity supports higher equilibrium stock. Verdict: ROBUST.

Test 3: Refresh. The reduced form produces slope $\partial L/\partial \mu = +8.38$. The Bayesian-with-noise rule produces slope $+1.32$. The threshold-cascade rule produces slope $+1.40$. All three rules produce positive slope, confirming the framework's prediction that higher refresh capacity supports higher equilibrium stock. Verdict: ROBUST.

Test 4: Rotation. The reduced form produces slope $\partial L/\partial \eta = +1.35$. The Bayesian-with-noise rule produces slope $+0.16$. The threshold-cascade rule produces slope $+0.26$. All three rules produce positive slope, confirming the framework's prediction that higher rotation supports higher equilibrium stock. Verdict: ROBUST.

7.2.3 Interpretation

All four comparative statics survive both alternative aggregation rules. The framework's qualitative predictions are robust to the specific aggregation assumption. The audience-formation gap that the framework acknowledges as a scope condition turns out to hide micro-detail rather than first-order effects.

This vindicates the framework's scope choice. Treating audience properties as observable primitives rather than deriving them from microfoundations does not create implicit dependence on a specific class of audience-formation models. The framework's results apply across reasonable aggregation rules, including ones that differ substantially in shape and in their treatment of audience heterogeneity.

A useful sub-finding emerges from the simulation. The absolute equilibrium stock levels differ substantially across aggregation rules, but the comparative statics are consistent. The Bayesian and threshold-cascade rules sustain higher absolute stocks than the reduced form because their sharper transition shape causes substrates to face stronger marginal disincentive against extraction in the relevant range. The reduced form's gradual response permits more substrate extraction at any given stock level, producing lower equilibrium stocks. But the direction of change in response to design parameters is consistent across all three rules.

The audience-formation sensitivity check therefore establishes that the framework's prescriptive program does not depend on a specific class of audience-formation models. The framework's qualitative results survive across the aggregation rules tested, and there is no reason to expect them to fail under other reasonable rules. The scope condition is honest rather than evasive: the framework declines to derive audience properties from microfoundations and acknowledges this as a limitation,

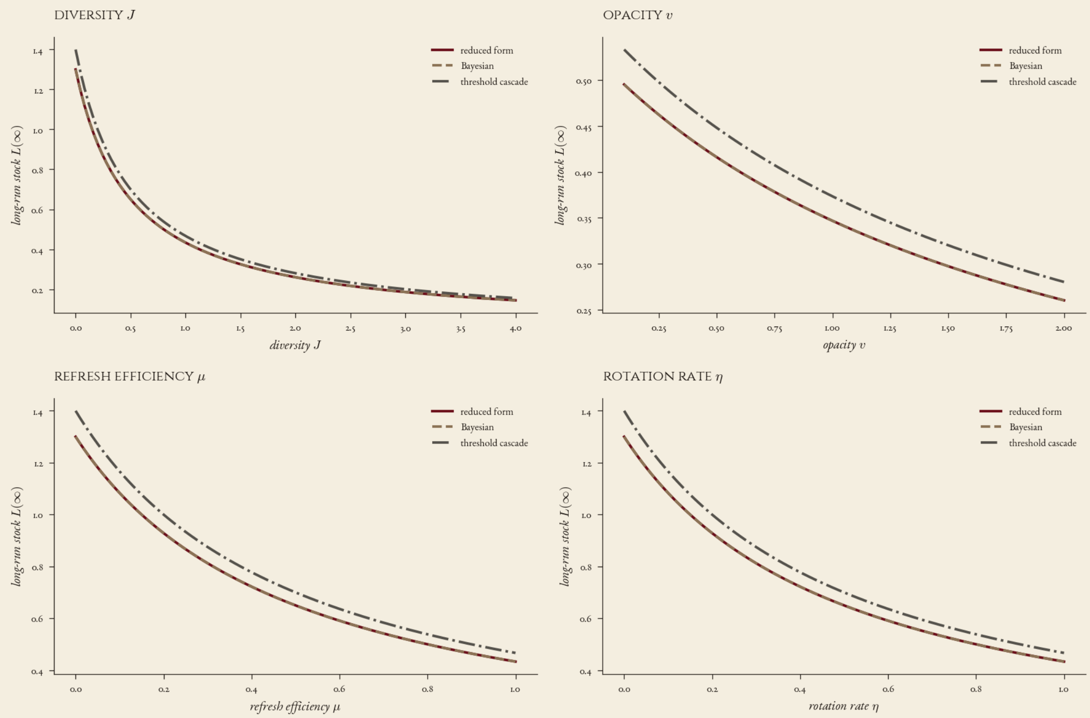

FIGURE 7.1. *Audience-formation sensitivity for the four comparative statics. Each panel shows long-run equilibrium stock* $L(\infty)$ *as a function of one design parameter (diversity* J*, opacity* v*, refresh efficiency* μ*, rotation rate* η*), under each of the three aggregation rules. The slopes have the same sign across all three rules in all four tests, confirming that the framework's qualitative comparative statics are robust to the audience-formation aggregation rule. Absolute stock levels differ across rules, but the directions of change are consistent.*

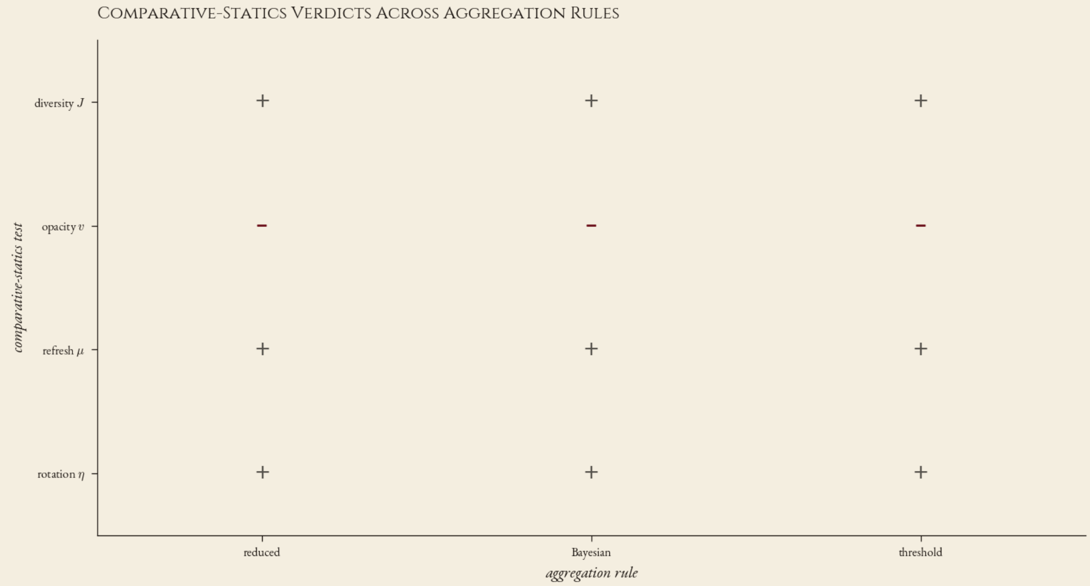

FIGURE 7.2. *Summary verdict on the audience-formation sensitivity check. All twelve cells (four comparative statics times three aggregation rules) PASS, with sign-consistent slopes across rules. The framework's qualitative predictions about diversity, detection, refresh, and rotation are robust to the audience-formation aggregation rule, supporting the scope choice to treat audience properties as observable primitives.*

but the limitation does not affect the framework's substantive predictions.

7.3 Remaining Limitations and Scope Conditions

The framework has substantive analytical reach but is not unlimited. Several limitations warrant explicit acknowledgment as scope conditions rather than implicit absences. Stating these upfront is part of the framework's methodological commitment to honesty about what it can and cannot establish.

7.3.1 The Audience-Formation Limitation

The framework treats audience properties as observable primitives rather than deriving them from microfoundational models. This is the limitation addressed by the sensitivity analysis above. The framework's qualitative results survive across reasonable aggregation rules, but the framework does not claim to provide microfoundations for audience formation. Readers seeking such microfoundations should consult the global games literature (Morris and Shin 2003, and successors) for the formal apparatus that the framework's reduced form approximates.

7.3.2 The Binary-Capture Simplification

The framework's Module 4 analysis treats constraint resource capture as a binary state: a resource is either captured or uncaptured. Real captures are typically gradual, with evaluation functions progressively biased rather than fully replaced. The simplification is deliberate and the framework's structural results extend to the continuous case (Φ_j^θ for $\theta \in [0, 1]$) at the cost of substantially longer proofs without changing the qualitative conclusions.

7.3.3 The Coverage Graph as Primitive

The framework's Module 4 analysis takes the coverage graph G_C as a primitive determined by institutional design. A complete treatment would derive the coverage graph from the substantive properties of each constraint resource (what actions it can evaluate, what evidence it requires) and the institutional design that connects them (jurisdictional authority, professional standards). The framework treats the graph as given and analyzes its topological properties, with prescriptions in Chapter 5 treating the graph as a design variable. Whether dense bidirectional structures can be implemented in practice is a substantive political question that the framework does not address.

7.3.4 The Reduced-Form Revolt Function

The framework adopts the reduced form $r(\tau, h) = \Phi((\tau - \tau_c)/\sqrt{v})$ for the population's revolt probability without deriving it from a global games or related microfoundational setup. The framework's substantive results follow from four qualitative properties of r (monotonicity in τ, τ_c, v, and h in the substantive range). Any defensible underlying model that produces a function with these properties is consistent with the framework's substantive results. The choice among such functions does not affect the framework's comparative statics, as the audience-formation sensitivity analysis confirms.

7.3.5 Single-Polity Analysis

The framework analyzes a single polity in isolation. International relations, inter-polity competition, and the dynamics of substrate behavior under exit options to other polities are not formalized within the framework. A multi-polity extension would model substrate behavior under exit competition, with citizens able to relocate to alternative polities and substrates competing for population on dimensions including extraction rates and constraint capacity. The framework's substantive results would change in identifiable ways under such an extension. Specifically, exit options would reduce the substrate's optimal extraction in equilibrium, since loss of population reduces sunk capital value. This is a known result in the literature on Tiebout competition and federalism. The framework's single-polity analysis is at its strongest where exit options are limited, which is true for most nation-state-scale polities under current conditions.

7.3.6 Empirical-Empirical Calibration

The framework's numerical results (the fifteen percent figure for Regime A, the sixty-five percent figure for Regime C, etc.) depend on specific calibration choices made for the simulation. The framework's qualitative results are calibration-independent, but the specific numbers are not. A more empirically grounded calibration would require historical data on substrate extraction, opacity, coordination cost, and renewal capacity across many polities, which is the empirical work of part two rather than part one.

7.3.7 *The Cases Await Their Demonstration*

The framework's contemporary cases (post-9/11 expansion of executive authority, post-2008 financial regulation, surveillance state development, and others discussed briefly in the introduction) are illustrations of the framework's analytical apparatus rather than empirical demonstrations. A demonstration would require systematic case-study analysis applying the framework's apparatus to specific historical episodes with full citation of evidence and engagement with the historiographical literatures. This work is deferred to part two.

7.3.8 *Prescription Implementation*

The framework's prescriptions identify what institutional features matter (resource diversity, capture detection, refresh capacity, capture rotation) without specifying how to implement them in any particular polity. The implementation question is substantively political and depends on the polity's existing institutional arrangements, political culture, and substantive context. The framework's contribution is identifying what to aim for; the implementation question is the work of part two.

7.4 The Framework's Analytical Standing

The three integrity tests of this chapter establish the framework's analytical standing at a specific level. The framework is internally consistent (all modules cohere, no formal results contradict each other, the notation is collision-free at source). The framework's qualitative predictions are robust to reasonable variation in modeling choices (the four key comparative statics survive across three audience-formation aggregation rules). The framework's limitations are stated explicitly as scope conditions rather than left implicit (the audience-formation gap, the binary-capture simplification, the coverage-graph-as-primitive treatment, the single-polity analysis, the calibration-dependent numerical results, the prescription implementation question).

This is what the framework can claim. What it cannot claim is unlimited analytical reach. The framework analyzes the dynamics of substrate authority over operational time within polities where substrate is necessary at scale; it does not analyze pre-substrate small-scale arrangements, multi-polity interactions, or the specific implementation of its prescriptions in any particular case. The scope conditions delimit the framework's claims and identify where additional analytical work would be required to extend the framework.

The methodological commitment to reporting limitations rather than concealing them is what gives the framework's substantive claims their force. A theory that claimed to derive everything from neutral premises and to support a flattering prescriptive program without limits would be unworthy of engagement on its own terms. The framework's willingness to identify its limits, and to subject itself to integrity checks that report what they find, is what makes the rest of its claims defensible. The reader is being asked to engage with a theory whose limits are visible, not with a theory whose limits are hidden.

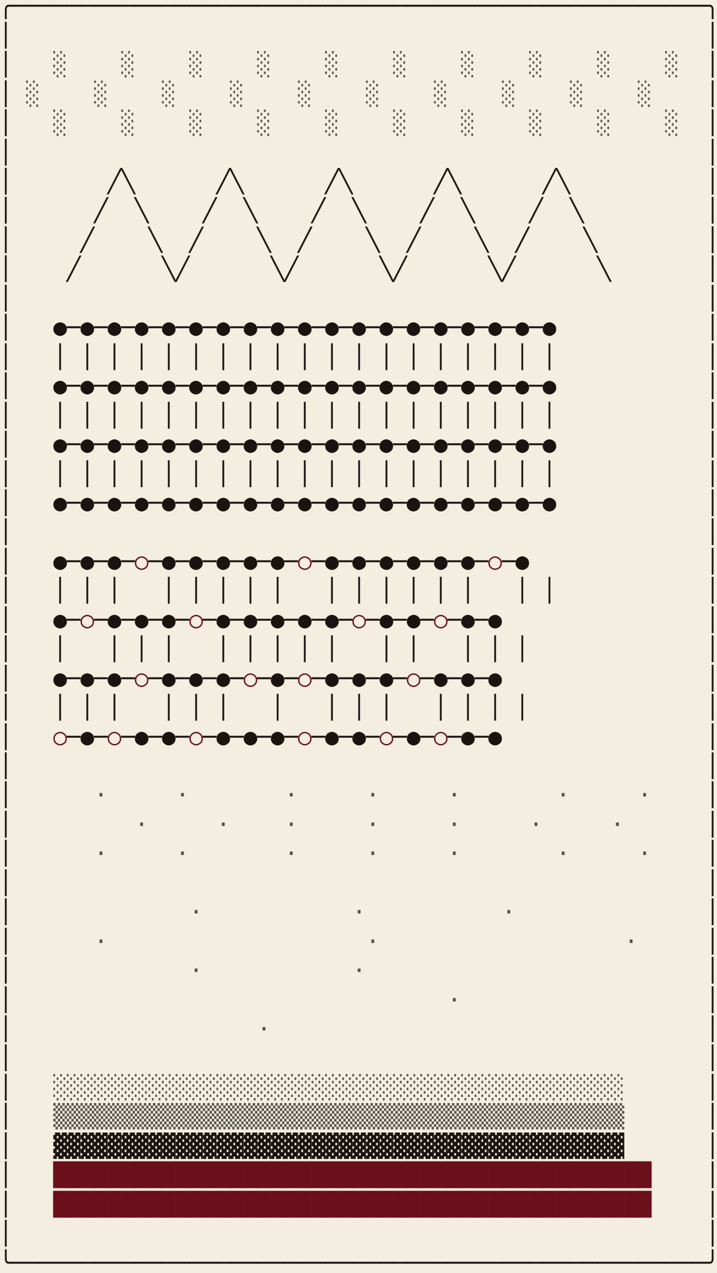

Conclusion

The framework presented across this volume answers a question that mainstream libertarian thought has largely declined to take up: what happens to a libertarian polity over operational time once the substrate that enables voluntary association at scale exists and begins to operate? The question is operational rather than founding. It concerns the dynamics of substrate behavior across years and decades, not the design of the substrate at its constitutive moment. The framework was built to provide analytical apparatus for this question, and the seven modules of analytical structure it has developed produce a coherent set of predictions about substrate trajectory under different institutional designs.

The headline predictions can be stated compactly. Voluntary association at scale requires substrate. Substrate is strategic and tends toward rent extraction. The constraint apparatus surrounding substrate is itself subject to capture dynamics, with cascading depletion across the resource portfolio and topological vulnerability in the coverage graph. Renewal events are themselves subject to capture under common conditions, with failed renewals producing consolidation rather than reset. Renewal capacity is multidimensional and admits multiple meaningful thresholds, with strict inclusion among the threshold regions. Specific institutional designs (resource diversity, capture detection, refresh capacity, capture rotation) substantially slow but do not prevent the trajectory from operational equilibrium toward terminal capture. Joint implementation of all four prescriptions does not eliminate substrate extraction but stabilizes it at the Irreducible Floor: long-run extraction at approximately fifteen percent of the unbounded maximum, not at zero.

These predictions distinguish Renewal Libertarianism from steady-state libertarian alternatives. Anarcho-capitalism's denial of substrate-necessity is wrong on the framework's apparatus. Minarchism's expectation of substrate stability after a properly designed founding is wrong on the framework's apparatus. The framework agrees with both alternatives on the foundational commitment to voluntary association as the ideal state for interactions, and parts company with both on what follows from that commitment given that substrate is necessary at scale and strategic over time.

The framework's substantive program, the elaboration of Renewal Libertarianism as a political position with specific policy recommendations, judgments about contemporary cases, and engagement with the substantive arguments of competing positions, is the subject of part two. Part one has established what the framework analytically shows. Part two will develop what to do about it.

The substantive program is organized around a single unifying claim. Renewal Libertarianism, as a political position, is the pursuit of **strategicity-handicap**: the cumulative reduction of substrate's capacity to optimize against population, sustained across operational time and punctuated by renewal events. This is a third position distinct from the two libertarian utopias the framework rules out. Anarcho-capitalism aims at substrate-elimination, which the substrate-necessity result rules out at scale. Minarchism aims at substrate-stability after a properly designed founding, which the Irreducible Floor rules out as a structural matter. Strategicity-handicap accepts substrate's existence and its necessity at scale; it targets substrate's capacity to optimize against population rather than substrate's existence or its quantitative behavior. The handicap operates through cumulative legitimacy costs of strategic action under what part two will name the **forcing-into-the-open principle**: well-designed institutional rules do not prevent substrate's strategic adaptation, which is structurally undefeatable, but force the adaptation to occur in publicly identifiable forms whose cumulative legitimacy cost handicaps substrate's capacity to strategize freely. Over generations, the institutional culture of substrate shifts as the cost-benefit ratio of strategic versus legitimate action moves against strategic action. This is the actual mechanism by which Western constitutional democracies became different in operation from their absolutist predecessors, and it is the mechanism the framework's apparatus characterizes.

The unifying claim has consequences for what part two delivers and for what it does not.

Part two delivers six specific contributions, each developed as part of the strategicity-handicap argument rather than as a stand-alone item.

The first contribution is the **forcing-into-the-open machinery**: the framework's four prescriptions translated into specific institutional recommendations whose cumulative effect is sustained strategicity-handicap accumulation. Resource diversity becomes the federalist redundancy and authority-type separation that prevents single-point capture. Capture detection becomes the rule structure (transparency requirements, mandatory disclosure, public default states with documented exceptions) that forces substrate adaptation into publicly visible form. Refresh capacity becomes Ostrom's design principles applied to legitimacy resources, with active support for civic participation through selective benefits and solidarity goods rather than mandatory participation. Capture rotation becomes term limits applied selectively, sortition for specific functions, procedural reset mechanisms, and federalism with mobility. Each is articulated as a contribution to the cumulative handicap target, with attention to the substrate-adaptation responses each forces and the auxiliary institutions each requires.

The second contribution is the **discrete-pulse component of strategicity-handicap accumulation**: the framework's cyclical structure implies that renewal events are the moments at which accumulated legitimacy costs come due all at once. Substrate that has been sustaining the marginal costs of forcing-into-the-open conditions across operational time encounters, at renewal events, the consolidated cost of every adaptation it has been forced to make. Part two develops the substantive recommendations for renewal frequency, structure, and constitutional convention mechanisms that maximize the strategicity-handicap effect of each renewal without triggering the renewal capture dynamics that part one characterizes. Renewal is the load-bearing political activity that distinguishes Renewal Libertarianism from competing libertarian positions; part two articulates this in detail.

The third contribution is the **empirical calibration of strategicity-handicap rates** through systematic case-study analysis of the contemporary cases part one has named at the level of illustration.

The post-9/11 expansion of executive authority, post-2008 financial regulation, the administrative state's long trajectory, and the surveillance state across multiple administrations are studied not just as instances of the pattern but as cases where strategicity-handicap accumulation can be measured. Part two extends the analysis to cases part one has not addressed and produces an empirical characterization of where strategicity-handicap has accumulated in the contemporary American configuration and where it has not.

The fourth contribution is the **sorting of competing libertarian and adjacent positions** by their relationship to strategicity-handicap. Part two engages Rothbard's *For a New Liberty* and *The Ethics of Liberty* and Hoppe's *Democracy: The God That Failed* on the substrate-elimination target, which the framework rules out. It engages David Friedman's *The Machinery of Freedom* on the market-emergence-of-constraint-resources question, which the framework engages directly. It engages the contemporary minarchist literature on the substrate-stability target, which the framework rules out. It engages Yarvin's neocameralist proposal on the strategicity-amplification counter (formal sovereignty without constraint apparatus, taking the framework's diagnosis as broadly correct but drawing the opposite prescriptive conclusion) and Land's accelerationist analysis on the strategicity-acceleration counter (the prescriptive question of whether to accelerate substrate's strategic dynamics toward a different terminal state). Each position is engaged on the substantive ground of what it does or does not do for strategicity-handicap, with the framework's apparatus as the analytical resource for the engagement.

The fifth contribution is the **contemporary American calibration**: which institutional reforms in which order would produce the largest strategicity-handicap accumulation given the contemporary American configuration's specific cascade-vulnerabilities and resource depletion patterns. This is where the implementation question gets specific. Part two identifies the high-priority reforms, the auxiliary institutions each requires, the substrate-adaptation responses each will provoke, and the rate at which strategicity-handicap is predicted to accumulate under sustained organizing for the reform program.

The sixth contribution is methodological. Organizing for strategicity-handicap requires knowing whether the organizing is producing effects, and the framework's comparative-statics analysis tells the libertarian what the long-run direction of effects is but not what intermediate observables would confirm being on track. Part two develops an **intermediate-observables methodology** for real-time monitoring of strategicity-handicap accumulation, drawing on the framework's apparatus to identify which observable patterns correspond to which underlying dynamics. The methodology is itself a contribution to the prescriptive program because the libertarian who is working for strategicity-handicap can know whether the work is producing effects.

What part two does not deliver, and why.

Part two does not deliver a precise formal characterization of strategicity-handicap as an analytical target. The framework's apparatus supports the direction of strategicity-handicap accumulation under sustained forcing-into-the-open conditions; it does not currently support a closed-form characterization of the precise institutional state at which strategicity-handicap is achieved. Part two is honest about this. The strategicity-handicap target is articulated as a qualitative goal whose conditions are characterized by the framework's existing apparatus, with the formalization work explicitly identified as analytical work that remains open. Future contributions to the framework's apparatus, by other analysts working in the framework's tradition or by analysts working outside it, may improve the formal characterization. Part two leaves room for those contributions and treats the current articulation as provisional rather than final.

Part two does not deliver a guaranteed timeline for strategicity-handicap accumulation. The framework's apparatus supports the direction; the rate at which strategicity-handicap accumulates under sustained reform depends on factors part two's empirical work calibrates rather than factors the apparatus delivers in advance. Part two estimates the rate for the contemporary American case based on the case-study work and provides intermediate observables for monitoring whether the actual rate is matching the estimated rate. The framework's claim is that strategicity-handicap is achievable on a human-lifetime horizon in favorable cases, with the contemporary American case admitting the favorable conditions if the reform program is pursued with sufficient organizing capacity. The framework cannot guarantee the favorable conditions will obtain.

Part two does not deliver a stable, final political position. Renewal Libertarianism, as part two articulates it, is the best current articulation of the libertarian commitment that the framework's apparatus supports. The framework's apparatus is not finished; it is open to extension in directions current articulations have not foreseen. Part two articulates Renewal Libertarianism with the explicit acknowledgment that successor articulations, drawing on framework extensions or alternative analytical apparatuses, may improve on the current position. The structural commitment is not that Renewal Libertarianism is the correct libertarian position but that it is the best current articulation given the framework's current apparatus, and that successor articulations more accurately tracking what the libertarian commitment is trying to characterize would be welcome and are the eventual goal. The framework's substrate-as-strategic-player diagnosis applies to the framework's own articulation; an ideology that aims at being supersedable is an ideology that has handicapped its own strategicity in the same way the framework's prescriptive program aims at handicapping substrate's strategicity.

A framework that predicted easy libertarian success would be untrustworthy on its face. A framework that predicted impossible libertarian failure would be uninteresting because it would offer no guidance. Renewal Libertarianism is positioned at the analytical place between these alternatives, taking the difficulty of the political world seriously and identifying what is achievable within it. The Irreducible Floor locates the achievable quantitative position; strategicity-handicap names the achievable qualitative target. The combination is the framework's substantive contribution. It is more sober than either of the libertarian utopias the framework rules out and more demanding than the libertarian movement's traditional anti-state self-conception. It asks the libertarian to organize across decades for an institutional culture transformation that is achievable on a human-lifetime horizon in favorable cases, with the auxiliary commitment that the work itself, and the framework that supports it, are provisional contributions to a larger project that successor analyses will continue.

What follows in part two is the substantive program that this analytical foundation supports.

Bibliography

Acemoglu, Daron, and James A. Robinson. 2006. *Economic Origins of Dictatorship and Democracy*. Cambridge: Cambridge University Press.

Acemoglu, Daron, and James A. Robinson. 2012. *Why Nations Fail: The Origins of Power, Prosperity, and Poverty*. New York: Crown Publishers.

Ackerman, Bruce. 1991. *We the People, Volume 1: Foundations*. Cambridge: Harvard University Press.

Buchanan, James M., and Gordon Tullock. 1962. *The Calculus of Consent: Logical Foundations of Constitutional Democracy*. Ann Arbor: University of Michigan Press.

Chenoweth, Erica, and Maria J. Stephan. 2011. *Why Civil Resistance Works: The Strategic Logic of Nonviolent Conflict*. New York: Columbia University Press.

Friedman, David. 1973. *The Machinery of Freedom: Guide to a Radical Capitalism*. La Salle, Illinois: Open Court. Third edition, CreateSpace, 2014.

Goldstone, Jack A. 1991. *Revolution and Rebellion in the Early Modern World*. Berkeley: University of California Press.

Hamilton, Alexander, James Madison, and John Jay. 1788. *The Federalist Papers*. Reprint, edited by Clinton Rossiter, New York: Mentor, 1961.

Hayek, Friedrich A. 1945. "The Use of Knowledge in Society." *American Economic Review* 35 (4): 519–530.

Hayek, Friedrich A. 1976. *Law, Legislation and Liberty, Volume 2: The Mirage of Social Justice*. Chicago: University of Chicago Press.

Hoppe, Hans-Hermann. 2001. *Democracy: The God That Failed: The Economics and Politics of Monarchy, Democracy, and Natural Order*. New Brunswick: Transaction Publishers.

Land, Nick. 2013. *The Dark Enlightenment*. Self-published online, 2012–2013. Collected and circulated as a single document.

Levitsky, Steven, and Daniel Ziblatt. 2018. *How Democracies Die*. New York: Crown.

Lichbach, Mark Irving. 1995. *The Rebel's Dilemma*. Ann Arbor: University of Michigan Press.

Locke, John. 1690. *Two Treatises of Government*. Reprint, edited by Peter Laslett, Cambridge: Cambridge University Press, 1988.

Morris, Stephen, and Hyun Song Shin. 2003. "Global Games: Theory and Applications." In *Advances in Economics and Econometrics: Theory and Applications, Eighth World Congress, Volume 1*, edited by Mathias Dewatripont, Lars Peter Hansen, and Stephen J. Turnovsky, 56–114. Cambridge: Cambridge University Press.

Niskanen, William A. 1971. *Bureaucracy and Representative Government*. Chicago: Aldine-Atherton.

North, Douglass C., John Joseph Wallis, and Barry R. Weingast. 2009. *Violence and Social Orders: A Conceptual Framework for Interpreting Recorded Human History*. Cambridge: Cambridge University Press.

Olson, Mancur. 1965. *The Logic of Collective Action: Public Goods and the Theory of Groups*. Cambridge: Harvard University Press.

Ostrom, Elinor. 1990. *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge: Cambridge University Press.

Rothbard, Murray N. 1973. *For a New Liberty: The Libertarian Manifesto*. New York: Macmillan.

Rothbard, Murray N. 1982. *The Ethics of Liberty*. Atlantic Highlands, New Jersey: Humanities Press.

Skocpol, Theda. 1979. *States and Social Revolutions: A Comparative Analysis of France, Russia, and China*. Cambridge: Cambridge University Press.

Stringham, Edward P., editor. 2007. *Anarchy and the Law: The Political Economy of Choice*. New Brunswick: Transaction Publishers.

Stringham, Edward P. 2015. *Private Governance: Creating Order in Economic and Social Life*. Oxford: Oxford University Press.

Tiebout, Charles M. 1956. "A Pure Theory of Local Expenditures." *Journal of Political Economy* 64 (5): 416–424.

Tocqueville, Alexis de. 1840. *Democracy in America*. Reprint, translated by Harvey C. Mansfield and Delba Winthrop, Chicago: University of Chicago Press, 2000.

Tullock, Gordon. 1967. "The Welfare Costs of Tariffs, Monopolies, and Theft." *Western Economic Journal* 5 (3): 224–232.

Yarvin, Curtis (writing as Mencius Moldbug). 2009. "A Gentle Introduction to Unqualified Reservations." Essay sequence published on *Unqualified Reservations*, 2009. Subsequent essays collected through 2014; later writings published as *Gray Mirror* (Substack, 2020–).