When Frontiers Do Not Convert

Contestability, Correction, and Hysteresis in Socio-Technical Change

Nicolò Bellanca

Department of Economics and Management, University of Florence, Italy

nicolo.bellanca@unifi.it ORCID: 0000-0002-3809-3455

Working Paper

June 30, 2026

Abstract

Evolutionary and frontier-based accounts of technological change ask how innovation expands the space of feasible outcomes and how institutions select among them. For infrastructure-like technologies this paper places a prior question at the centre: under what conditions is frontier technological capacity converted into socially distributed generative capability? Building on, but departing from, evolutionary trajectory theory, utility-frontier approaches, and the recent literature on innovation directionality, the paper develops a Socio-Technical Trajectory Space distinguishing frontier capacity, contestability of deployment assets, correction capacity, generative capability, and welfare. The central claim is that frontier capacity does not imply generative capability: advance becomes socially generative only when deployment is sufficiently contestable and correctable, formally when effective corrective openness exceeds a threshold. Contestability is measured as the traffic-weighted openness of strategic assets, connected here to the political economy of platform infrastructure and to the contestability standard at the centre of recent digital-market regulation. The paper's unifying result is a theorem, not a taxonomy: a single scalar, the enclosure bias, is shown by a lattice-theoretic selection argument to generate three distinct barriers to conversion — a static threshold, a measurement distortion, and a hysteretic exit delay — as three projections of one monotone, memory-bearing object, yielding a falsifiable conversion-reversal threshold.

1. Introduction: The Conversion Problem

The political economy of technological change is often organized around a frontier metaphor. New technologies expand the set of feasible economic and social outcomes; institutions determine which possibilities are selected; political conflict concerns the distribution of gains and losses. The metaphor links innovation, distribution, and institutional choice within a common analytical geometry, and it has proved durable for good reason.

Evolutionary economics has long complicated this picture from within. Technologies are not points that shift an aggregate production function; they are trajectories embedded in paradigms, routines, and capabilities, selected and retained through processes that are historical, cumulative, and path-dependent (Nelson and Winter, 1982; Dosi, 1982; Freeman, 1987; Metcalfe, 1998; Saviotti and Pyka, 2004; Perez, 2002; Arthur, 1989; David, 1985). **The social value of**

an innovation is never given by the artifact alone. It depends on whether heterogeneous organizations can generate, select, and retain locally adapted variants of its use.

Acemoglu's Utility-Technology Possibilities Frontier gives the frontier approach one of its most sophisticated contemporary formulations, making technology endogenous to institutions and power: societies shape which technologies are developed, adopted, blocked, and politically sustained (Acemoglu, 2025; Acemoglu and Restrepo, 2018, 2019). A parallel and more recent literature asks a cognate question from a policy angle: who gives *direction* to innovation, and through what institutional means. Mission-oriented and transformative innovation policy treat directionality — steering innovation activity toward specific societal goals — as something that has to be actively constructed, not assumed (Schot and Steinmueller, 2018; Mazzucato, 2018; Weber and Rohracher, 2012). We share their premise that direction is not given by the frontier alone. But directionality asks *where* technology should go; this paper asks a question that is logically prior: *whether*, once a direction is chosen and frontier capacity expands, that capacity actually becomes a capability that society at large can use, adapt, and correct. A mission can be well chosen and still fail to convert, if the assets through which it is deployed are closed.

This paper starts from these insights but argues that, for contemporary infrastructure-like technologies, a prior question must be asked. The problem is not only which point on a feasibility frontier is selected, nor only which groups gain or lose, nor only which direction policy chooses. The prior problem is whether frontier technological capacity is converted into socially distributed generative capability at all. A society may experience rapid growth in frontier capacity while the channels through which that capacity is deployed become more concentrated, less contestable, and less correctable. **The frontier may advance while the social capacity to adapt, recombine, and revise it stagnates or deteriorates.**

This is the conversion problem, and it has an increasingly concrete referent. Over the past decade, a growing political-economy literature on digital platforms has documented exactly the structural shift this paper formalizes: infrastructures that once mediated open access reorganize around closed, proprietary control points — compute, cloud services, data pipelines, APIs, standards — through which an expanding technological frontier is routed (Srnicek, 2017; van Dijck, Poell and de Waal, 2018; Nieborg and Poell, 2018). Regulators have converged on the same vocabulary from a different direction: the European Union's Digital Markets Act explicitly targets "contestable and fair markets in the digital sector" and designates dominant platforms as gatekeepers (Regulation (EU) 2022/1925). Contestability, in other words, is not a term invented for this paper; it is the organizing concept of the most consequential digital-economy regulation enacted in the period covered by this literature. This paper gives it a formal microstructure and asks what governs its erosion.

The paper develops this question by distinguishing frontier capacity from social generative capability, and by showing that the gap between them is governed by two mediating dimensions — the contestability of the assets through which deployment is routed, and the system's capacity to correct deployment errors — combined into a single measure of effective corrective openness. Section 2 gives these constructs their formal definitions, together with a notation table that states, for each, what it is intended to measure and how it could in principle be observed; nothing introduced here is purely formal vocabulary without an empirical referent.

Contributions. The paper makes four connected contributions. First, it shifts the unit of analysis from utility combinations or innovation direction to socio-technical conversion architecture: not which point is selected on a frontier, nor which direction is chosen, but whether selected and well-directed frontier capacity is actually converted into usable capability. Second, it provides a measurable bridge from asset concentration to contestability, conditional rather than innocent: when closure attracts traffic toward assets more closed than the average asset from which that traffic is drawn, a contestability measured at fixed weights overstates true openness. Third, it shows that information moving upward does not substitute for correction moving outward. Fourth, and as the paper's central theorem, it shows that a single enclosure parameter governs three distinct barriers to conversion — a static threshold, a measurement distortion, and a hysteretic exit delay — not as three independent mechanisms but as three projections of one monotone object: a closure choice that is shown, by a lattice-theoretic selection result, to move predictably with the enclosure parameter, and whose dynamic trace acquires memory exactly when the reversibility of strategic assets is heterogeneous.

This last result is the paper's principal claim to a theorem rather than a taxonomy. The common-bias representation shows that, under conditions stated and checked individually, these are not a list but a structure, and that the structure is strong enough to make the paper's closing economic claim — that a society can become more technologically advanced and less capable at the same time — a corollary of the formal apparatus rather than a conclusion appended to it.

The argument has immediate relevance for artificial intelligence. AI is not merely a productivity shock. It is an infrastructure-like frontier whose deployment passes through compute, cloud services, foundation models, proprietary datasets, APIs, standards, energy systems, procurement routines, audit rights, and specialized skills — the same class of strategic assets that the platform political-economy literature has already begun to document for AI specifically. AI can therefore increase frontier capacity while lowering contestability, weakening correction, and failing to raise generative capability.

The paper proceeds as follows. Section 2 defines the core variables of the Socio-Technical Trajectory Space and states, for each, its proxy and intended observability. Section 3 develops the traffic-share measure of contestability and its separability benchmark, connecting it to the platform-power literature introduced above. Section 4 derives the conversion threshold. Section 5 extends it to complementary components and bottlenecks. Section 6 shows why surveillance is not correction. Section 7 develops the paper's central theorem: it microfounds the enclosure path through a dependency-stock dynamics, proves monotone selection by a lattice-theoretic argument, shows that the three barriers are isotone projections of the selected closure path, and shows that the dynamic barrier acquires hysteretic memory under a precise and testable dispersion condition. Section 8 applies the framework to AI as a second-order juncture. Section 9 concludes. Technical derivations are collected in Appendices A–F.

2. The Socio-Technical Trajectory Space: Variables and Separations

The framework is built to separate objects often compressed into a single notion of technological progress. A frontier may advance; deployment channels may concentrate; correction may weaken; social capability may stagnate; welfare may rise for some and fall for others. The space distinguishes five levels: frontier capacity Φ, contestability D, correction capacity C, social generative capability G_{cap}, and welfare W. A sixth variable, depletion B, is introduced dynamically.

Table 1: Notation and observability

Symbol	Name	Measures	Observable proxy in principle
Φ	Frontier capacity	Technological capability at the frontier	Frontier models; R&D expenditure; frontier patents; compute capacity
D	Contestability	Openness of deployment assets	Traffic shares by asset; corrected concentration indices; contractual clauses
C	Correction capacity	Capacity to translate errors into revision	Active complaint channels; audit rights; rollback and revision mechanisms
$\Omega(D, C)$	Effective corrective openness	Effective openness produced jointly by contestability and correction	Composite measure increasing in both D and C
G_{cap}	Generative capability	Distributed social capacity to absorb and adapt frontier capacity	Local diffusion; organizational adaptive capacity
W	Welfare	Welfare, kept analytically distinct from G_{cap}	Distributional and welfare indicators
B	Depletion	Cumulative erosion of competencies and capacities	Loss of public competence; organizational autonomy; adaptive routines

2.1 Frontier capacity is not social capability

Frontier technological capacity Φ denotes capabilities at the technological margin: frontier models, advanced scientific knowledge, computational resources, specialized engineering, technical infrastructures, frontier organizational routines. Φ is not a welfare or social-capability measure. It is an input into a possible conversion process. A society may possess or import frontier technologies without acquiring the distributed capacity to understand, adapt, maintain, contest, or redirect them. In evolutionary language, Φ expands the space of possible techniques; it does not generate the population of locally retained variants that constitute realized capability.

2.2 Social capability is not welfare

Social generative capability G_{cap} is the distributed capacity to absorb, implement, diffuse, recombine, and use frontier capacity productively. It is narrower than welfare and broader than productivity, including absorptive capacity, local adaptation, organizational learning, diffusion across heterogeneous contexts, recombination, public administrative capacity, worker and user competence, and resilience of routines. Welfare W is kept separate: a technology may raise G_{cap} without raising W if distributional effects are sufficiently adverse, or may raise W in the short run through windfall gains without raising G_{cap}.

2.3 Contestability is not correction

Contestability D refers to the openness of the assets through which Φ is deployed: infrastructure, software architectures, datasets, standards, cloud services, compute, APIs, intellectual property, procurement contracts, skills, certification systems, access points. It is not identical to formal competition: a market may contain multiple vendors while all relevant channels remain locked into proprietary standards. Correction capacity C is analytically distinct: it denotes the ability to transform errors, mismatches, harms, and peripheral signals into effective revision. A deployment architecture may be relatively open but weakly corrective. Correction requires more than voice; it requires that voice be connected to mechanisms capable of changing implementation. In selection terms, D governs whether variant deployments can exist at all; C governs whether the selection environment retains the successful ones and discards the failures.

2.4 Effective corrective openness

These distinctions form a sequence: $\Phi \to \Omega(D, C) \to G_{\text{cap}} \to W$. The arrow from Φ to G_{cap} is mediated by effective corrective openness $\Omega(D, C)$. In the benchmark, $\Omega(D, C) = DC$; more generally the paper requires only that Ω be increasing in both arguments, with partial derivatives of Ω with respect to D and to C strictly positive. The multiplicative benchmark says that openness of assets and capacity to correct are complements: if assets are fully closed, correction cannot travel through them; if correction capacity is absent, formal openness does not generate adaptive revision.

The distinction between D and C also clarifies why information is insufficient. A controller may observe local conditions, collect performance data, or solicit messages, yet if these signals do not alter authority over deployment they do not constitute correction. Information moves upward; correction requires that revision capacity move outward. This is why surveillance and message transmission do not substitute for effective corrective openness, a result formalized in Section 6 and Appendix D.

2.5 Direction and depletion

The framework is a theory of socio-technical routing, not of technology alone. The same frontier capacity may have different social effects depending on the assets through which it passes. An AI system deployed through open standards, auditable models, portable data, internal public capacity, worker voice, and rollback rights differs from one deployed through proprietary APIs, non-portable data, external cloud dependence, opaque models, and weak appeal. The frontier may be similar; the trajectory is not. Finally, the framework allows for depletion B when persistent non-conversion generates cumulative damage to skills, public capacity, organizational autonomy, ecological resilience, trust, or infrastructural independence. In evolutionary terms, depletion is the erosion of the capacity to generate variety: once local craft knowledge, public expertise, and adaptive routines are lost, the population of variants from which correction could draw is itself reduced.

3. Strategic Assets and Traffic-Weighted Contestability

This section makes contestability operational. The aim is to avoid a common weakness in socio-technical arguments: the claim that some infrastructures matter more than others without a disciplined way to aggregate them. The framework treats strategic assets as channels through which frontier capacity must pass; their importance is measured by the share of deployment traffic routed through them.

3.1 Strategic assets

Let $\Sigma_t = \{1, \dots, S_t\}$ denote the set of strategic deployment assets at time t. An asset is strategic when access to it significantly affects whether frontier capacity can be transformed into social generative capability. For AI these may include compute, cloud infrastructure, foundation models, proprietary datasets, APIs, interoperability standards, energy access, procurement systems, technical expertise, audit infrastructures, and domain-specific implementation routines. The asset space is historically variable, not fixed by theory in advance, itself an evolutionary feature, since general-purpose technologies redefine which assets are selected as strategic.

The notion of a strategic asset whose control confers power independent of formal ownership concentration has an established lineage in the recent political economy of digital infrastructure. Mann's (1984) concept of infrastructural power — the capacity to penetrate and organize social activity through control of logistical channels rather than through direct coercion — has been extended to platform infrastructures, where platforms acquire infrastructural roles while existing infrastructures are reorganized through platform logics (Plantin et al., 2018). Srnicek (2017) and van Dijck, Poell and de Waal (2018) document the broader transition from open digital architectures to closed, proprietary infrastructures of exactly the kind this section formalizes; Nieborg and Poell (2018) trace this transition specifically as a move from open models to data-based, enclosed infrastructures in cultural production — an empirical instance of the enclosure path formalized in Section 7. The framework below gives this transition a measurable structure rather than treating it as a qualitative trend.

3.2 Corrected concentration and openness

For each asset s, let $\kappa_s \in [0, 1]$ denote raw concentration. Raw concentration is not enough: a concentrated asset may remain contestable under access obligations, interoperability rules, public oversight, open standards, audit rights, or effective countervailing capacity; conversely an asset with several nominal providers may be practically closed if all rely on the same proprietary standard or switching costs are prohibitive. The relevant measure is corrected concentration $\tilde{\kappa}_s \in [0, 1]$, incorporating ownership concentration, contractual lock-in, switching costs, opacity, absence of audit, non-portability of data, weak interoperability, and limited external accountability (the contractual and asset-specificity dimensions of lock-in in the sense of Williamson, 1985). A value near zero indicates an open, contestable asset; a value near one an effectively closed one.

Let $h_s(\tilde{\kappa}_s)$ denote channel openness, with h_s decreasing, $h_s(0) = 1$, and $h_s(1) = 0$. The linear benchmark is $h_s(\tilde{\kappa}_s) = 1 - \tilde{\kappa}_s$; openness declines proportionally with corrected concentration. Nonlinear specifications may be appropriate when small increases in concentration generate large drops in access, or when regulation buffers concentration up to a point.

3.3 The traffic-share bridge

A naive aggregate would assign weights ω_s to assets, but the weights would appear arbitrary. The framework interprets the weights as traffic shares. Let $\pi_s \in [0, 1]$ denote the share of frontier deployment traffic routed through asset s. Under the linear benchmark this becomes $D = 1 - \sum_s \pi_s \tilde{\kappa}_s$. Appendix A derives the bridge formally (Lemma A.1). Its meaning is simple: if most deployment traffic passes through open assets, aggregate contestability is high; if most passes through concentrated and closed assets, it is low. This shifts the measurement problem from abstract institutional evaluation to deployment mapping.

3.4 The separability benchmark

The bridge rests on an assumption to be stated explicitly because Section 7 relaxes it. The measured value of D treats traffic shares π_s and corrected concentrations $\tilde{\kappa}_s$ as separable: D is evaluated at a fixed weight vector, as if the deployment route were independent of how closed each asset is. Formally, the static measure assumes the partial derivative $\partial\pi_s/\partial\tilde{\kappa}_s$ is zero. This is analytically convenient and empirically tractable. But it is substantively strong. If closure on an asset raises switching costs and channels future deployment through that same asset, then concentration attracts traffic and the partial derivative is positive. Under separability D is a well-defined index; once separability fails, a D computed at fixed weights overstates true openness under a condition made precise in Section 7. For the static argument we proceed under separability and flag it as the benchmark to be tested.

3.5 Measurement, exposure matrices, and the dependency ceiling

A practical measurement strategy uses an exposure matrix. Let $j = 1, \ldots, J$ index deployment projects, with importance weights q_j summing to one, and let $e_{js} \in [0, 1]$ measure project j's exposure to asset s. When exposures are row-normalized, a natural estimator of traffic shares is $\pi_s = \sum_j q_j e_{js}$ (Lemma A.2). Entries may be measured from procurement contracts, technical documentation, system architecture, or usage logs, or proxied through vendor dependence, cloud expenditure, API calls, data-hosting arrangements, audit clauses, switching costs, or expert coding.

A distinct case arises when a strategic asset lies outside the jurisdiction of the deploying actor. When compute, foundation models, or cloud services are controlled by entities beyond regulatory, contractual, or democratic reach, corrected concentration on that asset has a floor domestic policy cannot lower. Let $\bar{\kappa}_s^{\text{ext}}$ denote the asset-specific lower bound on corrected concentration implied by external control. Then contestability is bounded above (Lemma A.3): a polity with strong legal accountability and competent procurement may still exhibit low effective contestability if a large share of deployment traffic passes through externally controlled assets. This is the formal counterpart of technological sovereignty as a conversion problem.

3.6 Partial identification and robust classification

Traffic shares may be contested, hidden, or rapidly changing. The framework therefore allows for identified sets. Instead of a single value of D, the analyst constructs an interval $[D^-, D^+]$ over a feasible set of traffic-share vectors. Fixing correction capacity C, let D_C denote the required contestability level. The regime is robustly converting if D^- exceeds D_C, robustly non-converting if D^+ falls below D_C, and fragile if the interval straddles D_C (Lemma A.4). This gives robustness an operational meaning through a signed distance index ρ, without pretending traffic shares are fully observed. Uncertainty over traffic shares matters for classification only when the identified set straddles the required level.

4. The Conversion Threshold

Contestability is only one component of the conversion architecture. Frontier capacity becomes social generative capability only when deployment is both contestable and correctable. This section formalizes the basic conversion condition.

4.1 Open and gated deployment

Let Φ denote an increment of frontier capacity. In open and correctable channels, local actors possess context-specific information, retain authority over implementation, and can transform mismatch into revision; this is the channel in which adaptive variety is generated and retained. In gated channels, deployment is controlled by an actor with discretion over use; the controller may observe information, but local adaptation is constrained and deployment may be biased toward enclosure, dependency, extraction, or standardization, that is, toward the suppression of variety.

Let $\Omega(D, C) \in [0, 1]$ denote effective corrective openness: the share or intensity with which deployment occurs through channels both sufficiently contestable and sufficiently correctable. Let $\lambda > 0$ be the net yield under locally adapted deployment and $\psi > 0$ the net damage under gated, biased deployment. The first-order conversion equation is the openness-weighted average of the two channel yields:

$$\Delta G = \Phi \left[\Omega(D, C)\lambda - \{1 - \Omega(D, C)\}\psi\right]. \tag{1}$$

Conversion is positive if and only if

$$\Omega(D, C) > \Omega^* := \frac{\psi}{\lambda + \psi}. \tag{2}$$

The threshold is therefore a required level of effective corrective openness, not a property of frontier capacity alone.

5. Complementarity, Bottlenecks, and the Weighting of Assets

When generative capability depends on several complementary components — technical skills, worker voice, public administrative capacity — conversion is governed not by the average openness across components but, as complementarity strengthens, increasingly by the single weakest one. In the limit of strong complementarity, the relevant conversion threshold converges to the threshold of the bottleneck component alone (Proposition C.2, derived in full in Appendix C): a deployment architecture cannot compensate for a non-correctable bottleneck by being open elsewhere.

This bottleneck result has a sharper implication once the weighting of assets is made explicit. The traffic-share measure of Section 3 weights assets by deployment volume; an alternative is to weight by criticality — an asset's contribution to the loss of feasible deployment that would follow from its removal. The two coincide when assets are substitutable; they diverge when some are bottlenecks. Appendix C shows (Lemma C.3) that when the shift from flow weights to criticality weights moves weight toward assets of below-average openness, contestability measured by criticality is lower than contestability measured by volume. A system can therefore look contestable by volume of traffic and non-contestable by bottleneck criticality — the very metric used to assess contestability can mask the bottleneck that matters.

Two implications follow for what comes later in the paper. First, effective corrective openness must be assessed at the level of the components that actually constrain conversion; aggregate, volume-weighted openness can mislead. Second, policy cannot rely on expanding frontier capacity or average access; it must identify and open the specific bottleneck. The more complementary the technology, the less credible it is to infer social progress from frontier capacity alone.

6. Why Surveillance Is Not Correction

The conversion threshold depends on effective corrective openness, not on information alone. Many concentrated deployment architectures defend themselves by claiming to improve information: platforms collect usage data, cloud providers monitor performance, AI vendors receive feedback, agencies require reporting. These practices may increase what a central actor knows. They do not necessarily increase the capacity of affected actors to correct deployment. The difference is between observation and authority. Surveillance moves information upward; correction requires that information generate revisable implementation, enforceable challenge, adaptation rights, or redistribution of control over use.

The result can be stated sharply. Let a controller receive a message M about the local adaptation type and implement the posterior mean plus a systematic bias b reflecting retained discretion. Then the biased action reduces expected yield by exactly b^2 relative to the unbiased action under the same information (Proposition D.1). As the message becomes perfectly informative, posterior variance vanishes and yield rises, but it remains strictly below the open-channel benchmark whenever b is nonzero. The residual loss b^2 is the non-purchasable component: it cannot be removed by any improvement in information, only by transferring authority over use. Message-contingent incentives can raise the informativeness of M but cannot lower b, because b is a property of who decides, not of what is known (Proposition D.2). This is the structural counterpart, in the deployment setting, of the limits to strategic information transmission when sender and receiver interests diverge (Crawford and Sobel, 1982): a bias between the party that holds information and the party that decides caps what communication alone can achieve.

This is why monitoring and audit must be distinguished. Monitoring collects information about system behaviour; audit in the corrective sense requires rights to inspect, challenge, verify, suspend, or modify deployment. A proprietary AI system may provide performance metrics without exposing model behaviour, data provenance, failure modes, or contractual lock-in. Participatory procedures become corrective only when connected to enforceable consequences: redesign, rollback, procurement revision, portability, appeal, worker voice, public oversight, or changes in access rights. Otherwise they risk becoming absorptive devices that channel dissatisfaction into non-binding communication while leaving deployment unchanged.

For AI the point is acute. AI systems often expand surveillance capacity while correction declines: workers may be monitored without being able to contest algorithmic decisions; citizens may be scored without meaningful appeal; agencies may depend on vendors whose systems they cannot inspect; users may provide feedback that improves the provider's model while gaining no rights over data, portability, explanation, or exit. The architecture becomes informationally rich and correctively poor. The framework thus rejects a common substitution of surveillance for contestability, reporting for correction, and information for authority. Correction is not knowledge at the top; it is revisability in the structure of use.

7. The Common-Bias Representation: Monotone Selection with Memory

7.1 The enclosure path as a structural object, not a list of effects

Section 6 has shown why information moving upward — surveillance, monitoring, reporting — does not substitute for correction moving outward. The question this section answers is different and prior: what makes deployment architectures drift toward closure in the first place, and why do the consequences of that drift show up simultaneously as a higher conversion threshold, a distorted measurement of contestability, and a delayed exit from non-conversion once it sets in.

The earlier draft of this argument presented these as three separate comparative-statics results, linked informally by the observation that they all responded to the same bias parameter. That presentation invited a fair objection: a list of derivatives that happen to share a sign is not, by itself, a theory of why they share it. This section replaces the list with a structure. We show that there is a single object — the closure choice made by an actor who benefits from a given level of enclosure — whose behaviour as the enclosure parameter b rises is governed by one minimal condition, and that the three barriers are nothing other than three different ways of reading that one object: a static reading, a measurement reading, and a dynamic reading with memory.

7.2 A dependency-stock microfoundation for closure

We do not posit the link between biased deployment and asset closure directly. We derive it from an accumulation process. When an actor retains discretion to deploy a technology with bias b — favouring centrally convenient implementations over locally adapted ones — dependency on the asset through which deployment is routed accumulates as a stock, growing with the intensity of biased use and depreciating at a constant rate. In steady state, this dependency stock is an increasing function of b: more biased deployment generates more dependency. The practical closure of the asset — measured, for instance, by switching costs, opacity, non-portability, or reliance on vendor-controlled complements — is in turn an increasing function of the accumulated stock. The result is that closure and bias are connected through a microfoundation rather than assumed to move together: a society does not close access to an asset *because* a modeller has stipulated that it does, but because dependency rents make doing so privately attractive once enough use has accumulated.

This microfoundation matters for what follows because it is the only place where the rest of the section's results draw their economic content. Everything from here on is a consequence of two monotonicities — the dependency stock rises with b, and closure rises with the stock — composed with one selection argument.

7.3 Monotone selection of closure

Consider an actor who chooses how aggressively to close a strategic asset, trading off the dependency rents that closure captures against the cost of closing it. This is a standard discrete-or-continuous choice problem, except that the trade-off itself shifts with b: a higher bias raises the dependency stock, and a higher dependency stock raises the marginal rent from any given level of closure. This is precisely the condition known in monotone comparative statics as increasing differences: the marginal benefit of closing further is itself increasing in b.

Under this condition — and under no condition stronger than it — a classical result (Topkis, 1978; Milgrom and Shannon, 1994) guarantees that the actor's optimal closure choice cannot fall as b rises. We state this carefully: the result guarantees that the *set* of optimal choices moves up in the appropriate order, and that its largest and smallest elements are nondecreasing in b; if the optimal choice is unique at each b, that unique choice is nondecreasing. We work throughout with the largest such selection, which is the economically relevant one — the most enclosure-prone optimal response — and write it $a^+(b)$.

The importance of this step is that it requires nothing beyond the two monotonicities of Section 7.2. It does not require the linear contestability benchmark used elsewhere in the paper for tractability, and it does not require the Gaussian-quadratic benchmark used to derive the conversion threshold in closed form. The monotone selection result is therefore the load-bearing structural claim of the paper; the specific functional forms used elsewhere are conveniences for computing closed-form expressions, not assumptions the structural claim depends on.

7.4 The three barriers as readings of one closure path

With $a^+(b)$ established as nondecreasing, the three barriers to conversion introduced earlier in the paper can each be shown to be a nondecreasing function of $a^+(b)$, and therefore nondecreasing in b for the same underlying reason.

The static threshold. The conversion threshold $\Omega^*(b)$ of Section 4 depends on b through the gated-deployment damage term, which rises with closure. Since the threshold is an increasing function of that damage term, it is nondecreasing along the selected closure path.

The measurement distortion. The fixed-weight overstatement of contestability introduced in Section 3 arises precisely when traffic is endogenously pulled toward assets more closed than the source-weighted average. A rise in $a^+(b)$ is, by construction, a rise in the closure of the enclosure-prone asset relative to others; the routing condition of Section 3 then implies that the overstatement widens.

The exit requirement. The level of communicative credibility required before correction can resume is nondecreasing in inherited closure: more closure leaves a deeper deficit of trust to rebuild before peripheral actors are again willing to report mismatches truthfully.

None of these three derivations is new content. What is new is that they are no longer three separate sign computations performed on three separate functional forms. They are three monotone readings of the single object established in Section 7.3. This is the precise sense in which the static threshold, the measurement distortion, and — as Section 7.5 shows — the hysteretic exit delay are three projections of one enclosure path, and not three independent mechanisms that happen to move together.

7.5 Memory: why exit is not the mirror of entry

The static reading of Section 7.4 treats b as a parameter to be compared across regimes. But non-conversion, once it sets in, does not simply reverse when the conditions that produced it relax. Exit from a closed regime is observed to lag behind entry into it: credibility, once it has decayed, rebuilds more slowly than it eroded, and assets that became practically closed under pressure do not automatically reopen when the pressure lifts.

We formalize this by treating the realized level of closure — read off the selected path

$a^+(b(t))$ as b moves over time — as the input to an operator with memory, of the type studied in the mathematics of hysteresis (Preisach operators). The operator represents the closure response as the aggregate behaviour of a population of switch-like elements, each of which turns on when the input rises through some threshold and turns off only when the input later falls through a *lower* threshold. If — and only if — there is genuine dispersion between the on-threshold and the off-threshold across the relevant population of strategic assets, the aggregate response traces an upper branch on the way in and a lower branch on the way out, and these two branches do not coincide. The horizontal distance between them, measured in the enclosure parameter, is the hysteretic exit delay.

This condition is not a technical nicety; it carries the paper's substantive claim about *why* exit is hard. Threshold dispersion is the formal counterpart of heterogeneous asset reversibility: some strategic assets re-open easily once enclosure pressure falls — a vendor restores an open API, a procurement rule is relaxed — while others remain closed well past the point at which the original pressure has lifted, because switching costs, accumulated proprietary complements, or eroded trust do not unwind at the same rate at which they accumulated. Where reversibility is uniform across assets, the model predicts no hysteresis at all: the exit delay is exactly zero. The prediction is therefore falsifiable in a specific sense — it can fail, and it fails precisely when the empirical dispersion of asset reversibility is found to be negligible.

7.6 The common-bias representation

Putting the three pieces together: a single enclosure parameter b selects, through a minimal monotone-selection argument resting on a dependency-stock microfoundation, a closure response $a^+(b)$ that cannot fall as b rises. The static conversion threshold, the fixed-weight measurement distortion, and — once asset reversibility is heterogeneous — the hysteretic exit delay are three nondecreasing functionals of that one response.

The theorem does not add a fourth mechanism to the model already developed in Sections 3–6. It shows that the existing mechanisms are three projections of one monotone enclosure path: the static projection is the conversion threshold, the measurement projection is the fixed-weight overstatement of contestability, and the dynamic projection is the hysteretic exit delay generated by dispersed reversibility thresholds. The central economic claim of the paper — that a society can become more technologically advanced and less capable at the same time — is therefore a corollary of this structure, not a conclusion appended to it after the technical work is done.

7.7 The conversion-reversal threshold

The theorem yields one further, directly testable implication. Define the conversion margin as the gap between attained effective corrective openness and the conversion threshold. Because attained openness is nonincreasing along the selected closure path while the threshold is nondecreasing, this margin moves in one direction as b rises: if it starts positive and the bias is large enough, it eventually crosses zero. The crossing point, b^*, marks a regime boundary. Below it, additional frontier capacity converts into generative capability. Above it, additional frontier capacity — absent an offsetting fall in b or a sufficient rise in contestability or correction capacity — *reduces* generative capability rather than raising it.

Under the hysteresis result of Section 7.5, the threshold for *recovering* converting status after a regime has crossed b^* lies strictly below b^* itself: a society that has crossed into non-conversion must reduce the enclosure parameter further than the level at which it originally entered before correction resumes. This is the paper's principal falsifiable prediction: not a smooth trade-off between frontier advance and social capability, but a regime boundary with an asymmetric entry and exit, observable in principle wherever the dispersion of asset reversibility can be assessed.

7.8 Formal statement and proofs

The argument of 7.1–7.7 is given here in fully formal terms, with every assumption stated separately so that each may be checked on its own. Readers interested only in the economic content may skip to Section 8.

Setup

Let $b \in I = [0, \bar{b}]$ denote the enclosure parameter. Let $(\mathcal{A}, \leq)$ be a compact sublattice of $\mathbb{R}^m$, ordered componentwise. Elements $a \in \mathcal{A}$ are vectors of practical closure actions on strategic assets. The controller's reduced-form payoff is

$$\Lambda(a, b) = R(a, b) - K(a), \tag{3}$$

where R collects dependency rents and K collects closure costs. The reduced-form closure correspondence is

$$A^*(b) = \arg\max_{a \in \mathcal{A}} \Lambda(a, b). \tag{4}$$

When $A^*(b)$ is not single-valued, write $a^+(b)$ and $a^-(b)$ for its greatest and least selections whenever they exist. The paper's economically relevant branch is the enclosure-prone branch $a^+(b)$; if the maximizer is unique, write it simply as $a^*(b)$.

The dependency-stock microfoundation is represented by

$$L_s^*(b) = \frac{\mu g(b)}{\delta}, \qquad \mu, \delta > 0, \tag{5}$$

for the enclosure-prone asset s, with dependency rents mediated by a map $\varphi(L_s^*(b))$.

Assumption 1 (Monotone dependency rents on the relevant interval)**.** On the interval I used in the theorem, g is nondecreasing and φ is nondecreasing on $L_s^*(I)$. In the differentiable case, $g'(b) \geq 0$ and $\varphi'(L) \geq 0$ on the relevant ranges.

Assumption 2 (Lattice regularity)**.** For every $b \in I$, $\Lambda(\cdot, b)$ is upper semicontinuous and supermodular on $\mathcal{A}$. The correspondence $A^*(b)$ is nonempty. Moreover, Λ has increasing differences in (a, b) on $\mathcal{A} \times I$.

Remark 1 (Why the interval matters). Assumption 1 is global only if φ is monotone on the whole calibrated range. If dependency rents saturate and then decline, all results below hold only on subintervals on which $\varphi' \geq 0$. This is the first load-bearing empirical/numerical check.

Topkis selection

Lemma 2 (Increasing differences from the dependency stock)**.** *Under Assumption 1, the dependency-stock channel gives increasing differences in* (a_s, b)*. In the smooth benchmark,*

$$\frac{\partial^2 \Lambda}{\partial a_s\, \partial b} = \varphi'(L_s^*(b))\, \frac{\mu g'(b)}{\delta} \geq 0. \tag{6}$$

This sign does not use the Gaussian-quadratic conversion benchmark and does not use the linear traffic-weighted-contestability benchmark.

Proof. The stock $L_s^*(b)$ is nondecreasing in b by (5) and Assumption 1. Since φ is nondecreasing on the relevant stock range, the marginal rent from raising the closure action on asset s is nondecreasing in b. This is exactly increasing differences. When the objects are differentiable, the chain rule gives (6). Its sign is fixed by $\varphi' \geq 0$, $g' \geq 0$, and $\mu/\delta > 0$. $\square$

Proposition 3 (Monotone enclosure branch)**.** *Under Assumption 2,* $A^*(b)$ *is increasing in the strong set order. In particular, the greatest and least selections* $a^+(b)$ *and* $a^-(b)$ *are nondecreasing in* b*. If the maximizer is unique, the unique selection* $a^*(b)$ *is nondecreasing in* b*.*

Proof. This is Topkis's monotonicity theorem. Supermodularity of $\Lambda(\cdot, b)$ and increasing differences in (a, b) imply that the argmax correspondence is increasing in the strong set order. Hence its extremal selections are isotone. If $A^*(b)$ is single-valued, the unique selection is the extremal selection and is therefore isotone. $\square$

Remark 4 (Correction relative to the earlier draft). Topkis does not imply that an arbitrary selection from a multi-valued argmax is monotone. It implies monotonicity of the correspondence in the strong set order and of the extremal selections; arbitrary selections require either uniqueness or an explicit monotone-selection convention. The theorem below therefore uses the greatest enclosure-prone branch $a^+(b)$, or the unique $a^*(b)$ when uniqueness is available.

The three barriers as monotone projections

Let $a(b)$ denote either the unique selection $a^*(b)$ or the greatest selection $a^+(b)$. Define three barriers:

$$T(b) := \Omega^*(b) = \frac{\psi(a(b), b)}{\lambda + \psi(a(b), b)}, \tag{7}$$

$$M(b) := D^{fw}(b) - D^{true}(b), \tag{8}$$

$$H(b) := \text{the exit delay generated by the memory operator in Section 7.5.} \tag{9}$$

Here $\lambda > 0$ is the open-channel yield and $\psi > 0$ is gated-deployment damage. The Gaussian-quadratic benchmark is the special case $\psi(b) = b^2 + \tau_\theta^{-1} - \bar{y}$, or equivalently the normalization in which the closure/bias branch itself is indexed by b.

Assumption 3 (Monotone barrier maps)**.** On I:

(i) $\psi(a, b)$ is nondecreasing in a and in b;

(ii) the fixed-weight measurement gap is nondecreasing in practical closure: closure attracts traffic toward assets at least as closed as the source-weighted average from which that traffic is drawn;

(iii) communication/corrective effectiveness is nonincreasing in practical closure, so the credibility level required for exit is nondecreasing in closure.

Lemma 5 (Isotonicity of the three projections)**.** *Under Assumptions 2 and 3, $T(b)$ and $M(b)$ are nondecreasing in b. The dynamic barrier $H(b)$ is nondecreasing in the Preisach sense made precise in Proposition 7: stronger closure raises the memory-state from which exit must be unwound.*

Proof. By Proposition 3, $a(b)$ is nondecreasing. For T, the map $x \mapsto x/(\lambda + x)$ is increasing on $x > 0$, and $\psi(a(b), b)$ is nondecreasing by Assumption 3; hence T is nondecreasing. For M, the routing condition in Assumption 3 says exactly that an increase in closure raises the fixed-weight overstatement: traffic is endogenously pulled toward assets more closed than the source average. Thus M is nondecreasing. For H, practical closure is the input to the memory operator; an increase in the input switches weakly more relays into the closed state and cannot reduce the inherited memory state. The strict exit delay is proved in Proposition 7 under the activated-dispersion condition. □

Preisach memory and the exit delay

Let $u(t) = m(a(b(t))) \in [\underline{u}, \bar{u}]$ be a scalar closure index, with m isotone. The scalar reduction is not substantive: it is the index read by the memory operator. Multiple assets can be represented by integrating asset-specific Preisach operators.

Definition 6 (Relay and Preisach closure operator)**.** For thresholds $\beta < \alpha$, let $\gamma_{\alpha\beta}[u](t) \in \{0, 1\}$ be the relay that switches from 0 to 1 when $u(t)$ rises through α and switches from 1 to 0 when $u(t)$ falls through β. Let ν be a finite positive measure on the Preisach half-plane

$$\Delta = \{(\alpha, \beta) : \underline{u} \le \beta < \alpha \le \bar{u}\}.$$

The effective closure output is

$$\mathcal{E}(t) = \mathcal{P}_\nu[u](t) := \int_\Delta \gamma_{\alpha\beta}[u](t)\, \nu(d\alpha, d\beta). \tag{10}$$

For a monotone rise from low closure to a peak u^+ and a subsequent monotone fall, the ascending and descending branches are

$$\mathcal{E}^\uparrow(u) = \nu\{(\alpha, \beta) \in \Delta : \alpha \le u\}, \tag{11}$$

$$\mathcal{E}^\downarrow(u; u^+) = \nu\{(\alpha, \beta) \in \Delta : \alpha \le u^+,\ \beta \le u\}. \tag{12}$$

Hence, for $u \le u^+$,

$$\mathcal{E}^\downarrow(u; u^+) - \mathcal{E}^\uparrow(u) = \nu\{(\alpha, \beta) : \beta \le u < \alpha \le u^+\} \ge 0. \tag{13}$$

Assumption 4 (Activated threshold dispersion)**.** For the peak u^+ and for the output level used to define exit, there exists a set of positive ν-measure with $\beta \le u < \alpha \le u^+$. Equivalently, the reversal actually activates off-diagonal relays whose switching-down thresholds lie below their switching-up thresholds in the relevant input range.

Proposition 7 (Endogenous exit delay)**.** *Assume $u(t)$ is piecewise monotone and consider a path that rises to u^+ and then falls. Then:*

(i) the Preisach output has the wiping-out property along piecewise monotone input histories;

(ii) the descending branch weakly dominates the ascending branch, with branch gap given by (13);

(iii) the branch gap is strictly positive at an input u if and only if the set $\{(\alpha, \beta) : \beta \le u < \alpha \le u^+\}$ has positive ν-measure;

(iv) if exit is defined as the first downward crossing of an output level $\bar{E}$ that lies between the two branches, then the exit input u_{out} is strictly below the corresponding entry input u_{in}. In input units, $h_u := u_{in} - u_{out} > 0$; in calendar time,

$$\tau_{exit} = \int_{u_{out}}^{u_{in}} \frac{du}{|\dot{u}_{down}(u)|} > 0, \tag{14}$$

whenever the downward speed is finite and positive on the interval.

Conversely, if ν is concentrated on the diagonal $\alpha = \beta$ within the activated range, then (13) vanishes for every u, the loop collapses, and the Preisach contribution to the exit delay is zero.

Proof. For piecewise monotone inputs, the Preisach operator has finite reversal memory: when the input reaches a new extremum, smaller nested reversal loops are erased. This is the wiping-out property. Along the first upward branch, a relay is on exactly when $\alpha \le u$, which gives (11). After the input has reached u^+ and is descending, a relay that was activated remains on until the input falls below its lower threshold β; this gives (12). Since $\beta < \alpha$, the ascending-on set is contained in the descending-on set, and subtracting the two sets yields (13). The gap is strictly positive precisely when the displayed straddling set has positive ν-measure. If an exit level $\bar{E}$ lies between the branches, the descending path reaches $\bar{E}$ only at an input strictly smaller than the ascending input that produced that level; hence $u_{out} < u_{in}$. Formula (14) is the change of variables from input distance to calendar time. If all activated thresholds lie on the diagonal, no relay has different switching-on and switching-off thresholds; there is no branch gap and no Preisach exit delay. □

Remark 8 (Economic meaning). Threshold dispersion is not an adornment. It says that strategic assets differ in reversibility: some re-open when enclosure pressure falls, others remain closed until pressure falls much further. The hysteresis of non-conversion is therefore proportional to the dispersion of asset reversibility. This is the second load-bearing empirical check.

Cardinal theorem and falsifiable corollary

Theorem 9 (Common-bias representation: monotone selection with memory)**.** *Let Assumptions 1, 2, and 3 hold on an interval $I \subseteq [0, \bar{b}]$. Let $a(b)$ be the unique closure selection, or the greatest enclosure-prone selection, from $A^*(b)$. Then the bias-induced components of the three non-conversion barriers are comonotone on I:*

(i) the selected closure branch $a(b)$ is nondecreasing in b;

(ii) the static conversion threshold $T(b)$ is nondecreasing in b;

(iii) the fixed-weight measurement overstatement $M(b)$ is nondecreasing in b;

(iv) if, in addition, Assumption 4 holds for the reversal path, the dynamic exit barrier generated by the Preisach operator is strictly positive and nondecreasing in the inherited closure state.

Thus the static threshold, the measurement distortion, and the exit delay are not three independent mechanisms. They are the static, measurement, and dynamic projections of one monotone closure branch, with the dynamic projection becoming hysteretic exactly when the closure branch is processed through a dispersed switching-threshold distribution. In this precise sense, the final economic claim of the paper follows as a corollary: frontier capacity can rise while social generative capability falls because the frontier axis and the conversion-architecture axis are distinct, and the latter is moved by the enclosure parameter.

Proof. Part (i) is Proposition 3, whose load-bearing condition is the increasing-differences result in Lemma 2. Parts (ii) and (iii) are Lemma 5. Part (iv) is Proposition 7. The interpretive conclusion follows because the same parameter b selects a weakly higher closure branch; all three barriers are monotone readings of that branch, and the dynamic reading adds memory exactly when activated thresholds are dispersed. □

Corollary 10 (Conversion-reversal threshold)**.** *Let attained effective corrective openness be $\Omega(b) = \Omega(D(b), C(b))$. Suppose Ω is continuous and nonincreasing on I, while T is continuous and nondecreasing. If there exist $b_0, b_1 \in I$ with $\Omega(b_0) > T(b_0)$ and $\Omega(b_1) < T(b_1)$, then*

$$b^* := \inf\{b \in I : \Omega(b) \leq T(b)\} \tag{15}$$

exists. For every $b > b^*$, marginal frontier expansion reduces social generative capability unless it is accompanied by a reduction in b or by an increase in D or C large enough to restore $\Omega(b) > T(b)$. If the Preisach dispersion condition holds, the recovery threshold after a reversal path is strictly lower than the entry threshold: $b_{out} < b^*$.

Proof. Continuity and the two sign conditions imply a crossing of $\Omega(b) - T(b)$ on I. Since Ω is nonincreasing and T is nondecreasing, the non-conversion set $\{b : \Omega(b) \leq T(b)\}$ is an upper interval, so b^* is the unique boundary up to flat segments. The marginal sign of frontier expansion is proportional to

$$\Omega(b)\lambda - (1 - \Omega(b))\psi(b) = (\lambda + \psi(b))(\Omega(b) - T(b)),$$

which is negative whenever $\Omega(b) < T(b)$. The strict inequality $b_{out} < b^*$ follows from Proposition 7: after the upward crossing, the downward branch remains more closed at the same input, so recovery requires a lower input than the entry crossing. □

---

8. Artificial Intelligence as a Second-Order Socio-Technical Juncture

AI makes the conversion problem visible. Much debate treats AI as a frontier technology whose effects depend on productivity, automation, task creation, displacement, market power, or regulation. These questions matter but do not exhaust the problem. AI is an infrastructure-like technology whose deployment passes through the strategic assets catalogued in Section 3:

compute, foundation models, proprietary datasets, audit infrastructures, and the standards and procurement routines that connect them. These assets determine whether frontier capacity becomes distributed capability or dependency-producing deployment.

AI may sharply increase Φ while the increase fails to convert. If AI deployment is routed through closed cloud infrastructures, proprietary APIs, non-portable data, opaque models, weak audit rights, vendor-controlled updates, and limited internal expertise, effective corrective openness may fall even as frontier capacity rises. The critical configuration is Φ rising, D and C falling, and G_{cap} uncertain or falling. This is not paradoxical: it is what the conversion model predicts. The social system gains access to more advanced operations but loses the capacity to adapt, inspect, challenge, repair, or redirect them. AI then becomes an externally mediated service dependency rather than a distributed capability.

8.1 A stylized operationalization

Consider a public agency adopting AI for case management, fraud detection, document processing, welfare allocation, or tax compliance. The conversion variables can be mapped to observable proxies. The following table is offered as an illustration of the framework's variables, not as an empirical estimate; the measurement of traffic shares and corrected concentration in such a setting is the natural first application, left to dedicated empirical work.

Table 2: Stylized AI operationalization

Strategic asset	Traffic proxy	Corrected-concentration indicator	Correction channel
Cloud	Share of deployment on external cloud	Lock-in, jurisdiction, reversibility	Audit, exit, procurement clauses
Foundation model	API dependence	Opacity, restricted model access	Evaluation rights, model documentation, rollback
Data	Portability and provenance coverage	Non-portability	Data governance, citizen appeal
Skills	Internal technical capacity	Dependence on vendor expertise	Public technical capacity, training

In such a setting officials may use a more advanced system while understanding less about how decisions are produced; citizens may receive faster decisions but have weaker appeal rights; the administration may become more digital while losing the capacity to diagnose, modify, or replace the system. Φ rises while $\Omega(D, C)$ may fall. A similar pattern can occur in firms, where AI tools raise productivity in some tasks while shifting knowledge from workers and internal teams to external vendors, converting local expertise into training data controlled by the provider.

8.2 The juncture: changing the asset space

AI should be understood as a second-order juncture. A first-order shock changes outcomes within a relatively stable asset space; a second-order juncture changes the asset space itself, redefining which assets are strategic and how deployment traffic flows across them. This is the sense in which AI is a general-purpose technology: like earlier engines of growth it pervades deployment and reshapes complementary assets across sectors (Bresnahan and Trajtenberg, 1995; Helpman, 1998), but here the pervasiveness operates on the very set of channels through which correction must travel. Compute becomes more central, cloud dependence increases, foundation models become gatekeeping infrastructures, data governance becomes a condition of capability, APIs become channels of access and control, audit infrastructure becomes essential, and energy and chip supply chains become constraints on sovereignty. Formally, AI changes not only Φ but the strategic asset set and the traffic shares.

The regulatory designation of dominant AI-adjacent platforms as "gatekeepers" under the Digital Markets Act (Regulation (EU) 2022/1925) is an institutional recognition, in real time, of exactly this second-order juncture: a determination that control over specific infrastructural chokepoints — not general market share — is what now requires contestability obligations.

Lemma 11 (Juncture-induced classification change)**.** *Fix correction capacity C and the required contestability level D_C. Suppose a second-order juncture adds a newly strategic, externally controlled asset s' with corrected concentration at or above its closure floor and positive traffic share. Then contestability after the juncture falls below the required level whenever the product of that asset's traffic share and corrected concentration exceeds one minus the required level minus the external traffic already carried by pre-existing assets, even if corrected concentrations on all pre-existing assets are unchanged.*

The proof applies the dependency-ceiling bound to the enlarged asset set (Appendix E). The lemma formalizes a simple but important point: technological change may reduce contestability not by closing old assets but by creating new bottlenecks. A system contestable before a general-purpose technology may become non-contestable after deployment traffic is rerouted through externally controlled infrastructures. This closes the circle between Section 3's dependency ceiling, Section 7's amplification, and the AI case: if the newly strategic asset attracts traffic because of dependency-producing closure, Lemma 8.1 and the monotone-selection result of Section 7 describe the same process at different levels, the asset space changing while traffic within the new space is endogenously pulled toward the bottleneck. The three components are thereby linked by a common formal result rather than by analogy alone.

8.3 Conversion-oriented AI policy

The framework reframes AI policy. Standard debates focus on innovation incentives, productivity, labour-market adjustment, safety, privacy, competition, and redistribution. These remain necessary but do not answer the conversion question. A conversion-oriented agenda asks different questions: are data portable, models auditable, APIs interoperable; can public agencies inspect and challenge vendor systems; do workers and users have channels to contest automated decisions; are there rollback rights; is there internal technical capacity to evaluate, modify, or exit systems; are procurement contracts designed against lock-in; are standards open enough for recombination; can affected communities transform observed errors into institutional revision? These concern $\Omega(D, C)$, not merely Φ.

This avoids two opposite simplifications. AI is not inherently destructive: under sufficiently contestable and correctable deployment it may raise generative capability by expanding access to expertise, improving public services, augmenting workers, accelerating discovery, and supporting organizational learning. Nor is it automatically emancipatory because it expands the frontier. The same frontier capacity becomes dependency-producing if deployed through closed and biased channels. AI capitalism need not be anti-innovation: concentrated actors may promote rapid frontier expansion precisely because expansion increases the strategic importance of assets they control. The characteristic figure is not the incumbent who blocks technology but the frontier gatekeeper who accelerates innovation while routing it through proprietary infrastructures. Power operates not by stopping the frontier but by controlling the conditions under which the frontier becomes, or fails to become, social capability.

9. Discussion and Conclusion: From Frontier Politics to Conversion Politics

Frontier technological capacity does not automatically become social generative capability. The central object is not the frontier alone but the architecture through which the frontier is converted. When deployment channels are contestable and correctable, and adaptive variety can be generated and retained, technological advance can expand distributed capability. When deployment is routed through closed, biased, and non-corrective assets, the same advance can generate dependency, depletion, and loss of social capacity.

The argument begins from a friendly critique of both evolutionary and frontier-based political economy, and engages a third, more recent interlocutor in the directionality literature. Evolutionary theory correctly insists that technologies are trajectories, that diffusion is local re-invention, and that the value of a technique depends on the generation and retention of variety. Utility-frontier models, including the UTPF, correctly insist that institutions, power, and technological direction must be analyzed together. The directionality literature correctly insists that the direction of innovation is not given but constructed. The framework connects all three: it asks whether the control structure of deployment assets permits the variety-generating, variety-retaining processes that evolutionary theory describes, treats the answer as prior to the distributive question frontier models pose, and shows that even a well-directed mission can fail to convert if the assets through which it is deployed are closed.

The formal results are deliberately disciplined. Frontier capacity increases generative capability only if effective corrective openness exceeds $\Omega^* = \psi/(\lambda+\psi)$. The traffic-share bridge gives the framework empirical discipline: contestability is the expected openness of the channels through which deployment passes, operationalizable through exposure matrices, procurement data, dependencies, contracts, interoperability conditions, audit rights, switching costs, and expert coding, with uncertainty represented through identified sets rather than hidden inside arbitrary weights. The bottleneck result shows that when complementarity is strong the relevant threshold is governed by the weakest component, and the flow-versus-criticality distinction shows that the very metric of contestability can mask that bottleneck. The surveillance result shows that information is not correction.

The unifying result is the common-bias representation proved in Section 7. A single enclosure parameter is shown, by a monotone-selection argument resting on a dependency-stock microfoundation, to govern a closure response that cannot fall as the parameter rises; the static conversion threshold, the fixed-weight measurement distortion, and — once asset reversibility is heterogeneous — the hysteretic exit delay are three nondecreasing readings of that one response, not three independent mechanisms. The result is stated conditionally and honestly throughout: the selection result is global only where the underlying rent function is monotone; the hysteretic component is strictly positive only where switching thresholds are genuinely dispersed; both conditions are testable rather than assumed. Within these limits, the result establishes that persistent non-conversion is not a list of mechanisms but one mechanism observed at three timescales: instantaneous conversion, measurement, and transition dynamics. The conversion-reversal threshold b* derived from the theorem gives the paper's principal falsifiable prediction: a regime boundary in the enclosure parameter, not a smooth trade-off, with an exit threshold that lies strictly below the entry threshold whenever asset reversibility is dispersed.

The policy implication is not anti-technological. The framework does not call for slowing frontier innovation as such. It calls for building architectures of conversion: open standards, interoperable systems, audit rights, data portability, regulated access to essential infrastructures, internal public capacity, procurement competence, worker and user voice, appeal procedures, rollback rights, and institutions capable of translating local errors into revision. Redistribution of gains remains important, but redistribution after deployment cannot substitute for contestability and correction inside deployment.

The framework has limits. The threshold model is local and stylized; it provides no full welfare theory; it does not claim that all technological concentration produces non-conversion; the dynamic mechanisms are conditional, and their global validity rests on monotonicity and dispersion conditions that must be checked rather than assumed. The hardest empirical task is measuring traffic shares and corrected concentration across strategic assets, and the amplification result shows this task is harder than a static estimate suggests, since the weights respond to the very closure they are meant to measure. These limits are also a research agenda: estimate deployment traffic shares in concrete domains, compare conversion architectures across sectors and countries, identify bottlenecks in generative capability, test whether low effective corrective openness predicts dependency, depletion, or weak diffusion despite high frontier access, and assess directly whether asset-reversibility thresholds are in fact dispersed in the domains where the model predicts hysteresis. Historical comparisons with electricity, railways, and computing could clarify when frontier expansion became socially generative and when it produced locked-in dependence.

The central conclusion is straightforward. The politics of technology is not only a politics of the frontier. It is a politics of conversion and correction. Technological progress cannot be inferred from the growth of frontier capacity alone. It depends on whether societies build the institutions, infrastructures, and rights through which frontier capacity becomes distributed, revisable, and socially generative, through which, in evolutionary terms, adaptive variety can be generated and retained. Where effective corrective openness is high, the frontier can become capability. Where it is low, societies may become more technologically advanced and less capable at the same time — not as a closing remark, but as a corollary of the formal result.

Appendices A–F

Appendices A–D reproduce, in compact but self-contained form, the technical derivations of Sections 2–6. Appendix E gives the proof of Lemma 8.1 (Section 8). Appendix F states the scope and limits of the theorem of Section 7.

Appendix A. Traffic-Share Contestability and Robust Classification

The proofs in Appendices A–F are stated compactly but with enough intermediate steps to make the derivations reconstructable without a separate supplement.

Lemma A.1 (Traffic-share contestability). Let aggregate contestability be the expected openness of the channel through which a unit of frontier deployment passes. With traffic shares π_s held fixed,

$$D = \sum_s \pi_s h_s(\tilde{\kappa}_s).$$

Under the linear benchmark $h_s(\tilde{\kappa}_s) = 1 - \tilde{\kappa}_s$, this is

$$D = 1 - \sum_s \pi_s \tilde{\kappa}_s.$$

Proof. A unit of deployment is routed through asset s with probability π_s; conditional on routing through s, channel openness is $h_s(\tilde{\kappa}_s)$. Aggregate contestability is the expectation of channel openness over the routing distribution, which is the stated sum. The linear case follows by substitution and the fact that the shares sum to one. The lemma converts weights from discretionary normative judgments into traffic shares, shifting measurement to deployment mapping. The shares are held fixed here; Appendix E relaxes this.

Lemma A.2 (Traffic shares from an exposure matrix). Let projects be indexed by j, with importance weights q_j summing to one, and let $e_{js} \in [0, 1]$ be the exposure of project j to asset s. If exposures are row-normalized, so that each project's exposures sum to one, then $\pi_s = \sum_j q_j e_{js}$ is a valid traffic-share vector.

Proof. Non-negativity is immediate. Summing over s and exchanging the order of summation, the inner sum over s of e_{js} equals one for each j by row-normalization, leaving the sum over j of q_j, which is one. The row-normalization condition states that each project's deployment is fully allocated across strategic assets, a condition often imposable by construction when exposures are shares of cost, API calls, data-hosting requirements, or cloud dependence.

Lemma A.3 (Dependency ceiling). Let an external subset of assets be controlled outside the jurisdiction of the deploying actor, with corrected concentration bounded below by a floor on each (a closure floor, which produces a ceiling on contestability). Under the linear benchmark, contestability is bounded above by one minus the traffic-weighted sum of those floors over the external assets. More generally, with h_s decreasing, contestability is at most the external-asset openness plus the total non-external traffic share.

Proof. For external assets openness is at most one minus the closure floor, since h_s is decreasing; for non-external assets openness is at most one. Substituting into the bridge of Lemma A.1 and using that the shares sum to one gives the bound. The closure floor is the formal counterpart of technological sovereignty: domestic openness on the remaining assets cannot lift contestability above the bound imposed by the traffic routed through externally controlled assets. It is not a claim that external dependence is always harmful, only that closed external assets carrying traffic impose an upper bound.

Lemma A.4 (Identified set and robust classification). Let the feasible set of traffic-share vectors be non-empty, compact, and convex. Then contestability is partially identified in a compact interval $[D^-, D^+]$ obtained by minimizing and maximizing the linear functional of Lemma A.1 over the feasible set. Fixing correction capacity C with a required contestability level D_C, the regime is robustly converting if D^- exceeds D_C, robustly non-converting if D^+ is below D_C, and fragile if the interval straddles D_C.

Proof. For fixed corrected concentrations and openness functions, the map from a traffic-share vector to contestability is a continuous linear functional. On a non-empty compact set it attains its extrema, so the bounds are well defined; convexity makes the image a compact interval, so every intermediate value is attained and the identified set is exactly the interval. Since effective corrective openness is strictly increasing in D at fixed C, conversion occurs exactly when D exceeds D_C, and the three classification cases follow. A robustness index is the signed distance from the identified interval to D_C, giving robustness an operational meaning without pretending traffic shares are fully observed.

Appendix B. Gaussian-Quadratic Conversion Threshold

Let each frontier innovation have adaptation requirement θ with Gaussian prior of precision τ_θ, and let a local actor observe a signal equal to θ plus Gaussian noise of precision τ_ε. Implementation chooses an action a; gross potential is $\bar{y}$ and yield is $\bar{y}$ minus the squared mismatch between a and θ.

Lemma B.1 (Gaussian-quadratic yields). Under open deployment the actor chooses the posterior mean, with expected squared error equal to the posterior variance, the reciprocal of the sum of the two precisions. The open-channel yield is therefore $\lambda = \bar{y} - 1/(\tau_\theta + \tau_\varepsilon)$. Under baseline gated deployment the controller imposes a biased action b without using the local signal; since the prior mean is zero, expected squared error is $b^2 + 1/\tau_\theta$, and the gated damage when gated deployment is destructive is $\psi = b^2 + 1/\tau_\theta - \bar{y}$.

Proof. The posterior-mean error equals the posterior variance for a Gaussian model; the open yield follows. For the gated case the expected squared deviation of a fixed action from a zero-mean prior of variance equal to the reciprocal of τ_θ is the bias squared plus that variance; subtracting from gross potential and negating gives ψ when the gated yield is negative. The open-channel loss is informational; the gated-channel loss is institutional, arising from imposed bias and the failure to use local adaptation information.

Proposition B.2 (Conditional generativity). The same frontier innovation is productive under open deployment and damaging under gated deployment if and only if gross potential lies in the contested band between $1/(\tau_\theta + \tau_\varepsilon)$ and $1/\tau_\theta + b^2$. The band is non-empty whenever the signal precision is positive.

Proof. Open deployment is productive exactly when λ is positive, that is, when gross potential exceeds the lower bound; gated deployment is damaging exactly when ψ is positive, that is, when gross potential is below the upper bound. The intersection is the band, non-empty because the lower bound is strictly below the reciprocal of τ_θ, which is at most the upper bound. The band widens as τ_θ falls: radical or context-sensitive technologies make local knowledge more valuable and imposed deployment more dangerous.

Proposition B.3 (Conversion threshold). Let a fraction $\Omega(D, C)$ of deployment pass through open or effectively corrected channels. Then expected conversion per unit of frontier capacity is the openness-weighted average of λ and minus ψ, which is positive if and only if $\Omega(D, C)$ exceeds $\Omega^* = \psi/(\lambda + \psi)$.

Proof. Expected conversion is $\Omega\lambda - (1-\Omega)\psi$. Setting this positive and rearranging, using that $\lambda + \psi$ is positive, gives the stated condition. The threshold is a required level of effective corrective openness, not of contestability; only when Ω equals D does it reduce to a condition on D.

Lemma B.4 (Threshold comparative statics). With λ and ψ positive, the threshold is increasing in the squared bias, with derivative $\lambda/(\lambda+\psi)^2$; decreasing in gross potential (which enters both λ and ψ, while their sum is independent of it, so the derivative is $-1/(\lambda+\psi)$); decreasing in signal precision; and decreasing in prior precision. The earlier claim that the sign in prior precision is indeterminate does not hold once λ and ψ are restricted to be positive: both channels through which prior precision enters lower the threshold, so the derivative is strictly negative.

Proof. Writing the threshold as $\psi/(\lambda+\psi)$ and differentiating, the derivative with respect to the squared bias is $\lambda/(\lambda+\psi)^2 > 0$. Gross potential enters only ψ, with derivative -1, giving $-1/(\lambda+\psi) < 0$. Signal precision enters only the open yield through the combined precision, lowering posterior variance and thus lowering the threshold. Prior precision lowers gated damage and raises the open-channel yield; with $\lambda, \psi > 0$, both channels lower the threshold, so the derivative is strictly negative. Greater enclosure bias thus raises the openness required for conversion; greater gross potential, higher signal precision, and higher prior precision lower it.

Appendix C. CES Aggregation, Bottlenecks, and Asset Weighting

Lemma C.1 (CES shadow weights and local threshold). Let generative capability be a CES aggregate of components with substitution parameter σ, each component evolving with its own yield and damage. Solving for a common openness level, the local threshold for non-negative change in aggregate capability is a weighted average of the component thresholds, with weights proportional to each component's marginal CES share divided by its level and scaled by its yield-plus-damage; these weights rise when a component is scarce relative to the others. If openness is component-specific rather than common, there is no single scalar threshold and the condition is the corresponding weighted inequality.

Proof. Differentiating the CES aggregate, the marginal contribution of each component is proportional to its marginal CES share divided by its level. Setting the first-order change in the aggregate to zero and solving for the common openness level that holds aggregate generativity constant yields the stated weighted threshold, with weights proportional to (marginal share / level) times yield-plus-damage. The threshold is not an ordinary average of component thresholds: already-weak components receive higher weight because the aggregator becomes increasingly sensitive to bottlenecks as complementarity rises, that is, as σ falls toward zero.

Proposition C.2 (Bottleneck convergence). If there is a unique weakest component, then as complementarity rises, that is, as $\sigma \to 0$, the local threshold converges to the threshold of that bottleneck component.

Proof. As complementarity rises (that is, as $\sigma \to 0$ in this parametrization), the ratio of any non-bottleneck shadow weight to the bottleneck shadow weight tends to zero, because the bottleneck has the smallest level raised to a growing power. The bottleneck weight tends to one and all others to zero, so numerator and denominator of the local threshold collapse to their bottleneck terms. Openness elsewhere cannot compensate for a non-correctable bottleneck when complementarity is strong.

Lemma C.3 (Criticality dominates at bottlenecks). Define the flow share of an asset as its share of deployment volume and its criticality share as its normalized loss in feasible deployment under removal. If the reweighting from flow to criticality shares has non-positive covariance with channel openness — equivalently, if the net weight shift $\sum_s (c_s - \varphi_s)(h_s - \bar{h})$ is at most zero, so that mass moves on net toward assets of below-average openness — then contestability evaluated with criticality shares is below contestability evaluated with flow shares, strictly so when the net shift toward below-average assets is strict.

Proof. The difference between criticality-weighted and flow-weighted contestability is the sum over assets of the weight change times openness. If the reweighting moves mass toward below-average-openness assets, the covariance between weight change and openness is non-positive, so the difference is non-positive, and strictly negative when positive mass moves toward an asset of strictly below-average openness. A system may thus look contestable when measured by volume and non-contestable when measured by bottleneck criticality; the choice of weighting encodes a substantive view of the deployment architecture.

Appendix D. Surveillance, Messages, and Biased Use-Discretion

Proposition D.1 (Irreducibility of retained bias). Let a controller receive a message about the local adaptation type and implement the posterior mean plus a systematic bias b. Then the biased action reduces expected yield by exactly b^2 relative to the unbiased action under the same information.

Proof. The implementation error decomposes into the posterior-mean error plus the bias. Squaring and taking expectations, the cross term vanishes because the posterior-mean error is orthogonal to any measurable shift, leaving the expected posterior variance plus the squared bias. The unbiased action incurs only the posterior variance, so the biased action yields exactly b^2 less under the same message. In the limit of full information posterior variance vanishes, yet the squared bias remains. The non-purchasable object is not information but corrective authority over use.

Proposition D.2 (Scope under transfers). Transfers may induce more informative reporting and, in the limit, full revelation, but if the controller retains biased use-discretion the yield remains lower by b^2 than under the unbiased action on the same information.

Proof. Transfers affect the informativeness of the message; in the limit posterior variance vanishes. But the bias is a property of the controller's implementation rule, not of the message. By Proposition D.1 the yield gap of b^2 persists for any information structure. Elicitation is not the same as institutional correction: information can be elicited, but correction of retained use-discretion cannot be purchased by revelation alone.

Appendix E. Juncture-Induced Classification Change

After the juncture the relevant asset set is the old set together with the new externally controlled asset. By the dependency-ceiling bound of Lemma A.3 applied to the enlarged set, contestability after the juncture is at most one minus the new asset's traffic share times its corrected-concentration ceiling, minus the external traffic of pre-existing assets. The pre-existing concentrations may be unchanged, but the new asset now carries positive traffic; if the product of its traffic share and corrected concentration is large enough, the bound falls below the required contestability level. Technological change can thus reclassify a system as non-converting by creating a new bottleneck rather than by closing old assets.

Appendix F. Scope and Limits of the Common-Bias Representation

The results of Section 7.8 are stated under conditions that are checked individually rather than assumed as a block. Three scope conditions deserve explicit statement, since each marks the boundary of the result's validity rather than a decorative caveat.

Global versus local monotone selection. Proposition 3 is global on the full interval $I = [0, \bar{b}]$ only if the dependency-rent map φ is nondecreasing on the entire calibrated range of the dependency stock. If rents saturate and then decline at high stock levels — a plausible pattern if extreme dependency triggers regulatory intervention or substitution — the monotone-selection result, and everything built on it, holds only on the maximal subinterval(s) on which $\varphi' \geq 0$. Outside that interval the model does not predict the closure response is monotone, and the three barriers need not be comonotone.

Dispersion is necessary, not decorative. The hysteretic exit delay of Proposition 7 is strictly positive if and only if the distribution of asset switching thresholds places positive mass off the diagonal $\alpha = \beta$ in the activated range – equivalently, if and only if strategic assets genuinely differ in their reversibility. This is an empirical claim, not a free parameter of the model: it predicts that hysteresis should be observed precisely where asset reversibility is heterogeneous, and should vanish where it is not. The model is falsified, in this specific dimension, if dispersion is found to be negligible in a domain where it predicts a sizeable exit delay.

Piecewise-monotone paths. The closed-form branch formulas for the hysteretic exit delay assume that the enclosure parameter, and the closure index it induces, moves in a piecewise monotone way over the relevant time window: it rises to a peak and then falls, rather than oscillating repeatedly. The Preisach operator remains well defined under oscillatory paths, but the relevant object for measuring the exit delay is then the loop integral over the realized path rather than the simple closed-form gap given above.

None of these conditions undermines the structural claim of Section 7: that the static threshold, the measurement distortion, and the hysteretic exit delay are projections of one monotone, memory-bearing object rather than three independent mechanisms. They specify precisely where that claim is global and where it must be restricted to a subinterval or made conditional on an empirically assessable property of the asset population.

References

Acemoglu D (2025) Nobel lecture: institutions, technology, and prosperity. *American Economic Review* 115(6):1709–1748. doi:10.1257/aer.115.6.1709.

Acemoglu D, Restrepo P (2018) The race between man and machine: implications of technology for growth, factor shares, and employment. *American Economic Review* 108(6):1488–1542. doi:10.1257/aer.20160696.

Acemoglu D, Restrepo P (2019) Automation and new tasks: how technology displaces and reinstates labor. *Journal of Economic Perspectives* 33(2):3–30. doi:10.1257/jep.33.2.3.

Arthur WB (1989) Competing technologies, increasing returns, and lock-in by historical events. *The Economic Journal* 99(394):116–131. doi:10.2307/2234208.

Bresnahan TF, Trajtenberg M (1995) General purpose technologies: "Engines of growth?" *Journal of Econometrics* 65(1):83–108. doi:10.1016/0304-4076(94)01598-T.

Crawford VP, Sobel J (1982) Strategic information transmission. *Econometrica* 50(6):1431–1451. doi:10.2307/1913390.

David PA (1985) Clio and the economics of QWERTY. *American Economic Review* 75(2):332–337.

Dosi G (1982) Technological paradigms and technological trajectories: a suggested interpretation of the determinants and directions of technical change. *Research Policy* 11(3):147–162. doi:10.1016/0048-7333(82)90016-6.

Freeman C (1987) *Technology Policy and Economic Performance: Lessons from Japan.* Pinter Publishers, London.

Helpman E (ed) (1998) *General Purpose Technologies and Economic Growth.* MIT Press, Cambridge, MA.

Mann M (1984) The autonomous power of the state: its origins, mechanisms and results. *European Journal of Sociology* 25(2):185–213. doi:10.1017/S0003975600004239.

Mazzucato M (2018) Mission-oriented innovation policies: challenges and opportunities. *Industrial and Corporate Change* 27(5):803–815. doi:10.1093/icc/dty034.

Metcalfe JS (1998) *Evolutionary Economics and Creative Destruction.* Routledge, London.

Milgrom P, Shannon C (1994) Monotone comparative statics. *Econometrica* 62(1):157–180.

Nelson RR, Winter SG (1982) *An Evolutionary Theory of Economic Change.* Belknap Press of Harvard University Press, Cambridge, MA.

Nieborg DB, Poell T (2018) The platformization of cultural production: theorizing the contingent cultural commodity. *New Media & Society* 20(11):4275–4292. doi:10.1177/1461444818769694.

Perez C (2002) *Technological Revolutions and Financial Capital: The Dynamics of Bubbles and Golden Ages.* Edward Elgar, Cheltenham.

Plantin J-C, Lagoze C, Edwards PN, Sandvig C (2018) Infrastructure studies meet platform studies in the age of Google and Facebook. *New Media & Society* 20(1):293–310. doi:10.1177/1461444816661553.

Regulation (EU) 2022/1925 of the European Parliament and of the Council of 14 September 2022 on contestable and fair markets in the digital sector and amending Directives (EU) 2019/1937 and (EU) 2020/1828 (Digital Markets Act). *Official Journal of the European Union* L 265:1–66.

Saviotti PP, Pyka A (2004) Economic development by the creation of new sectors. *Journal of Evolutionary Economics* 14(1):1–35. doi:10.1007/s00191-003-0179-3.

Schot J, Steinmueller WE (2018) Three frames for innovation policy: R&D, systems of innovation and transformative change. *Research Policy* 47(9):1554–1567. doi:10.1016/j.respol.2018.08.011.

Srnicek N (2017) *Platform Capitalism.* Polity Press, Cambridge.

Topkis DM (1978) Minimizing a submodular function on a lattice. *Operations Research* 26(2):305–321. doi:10.1287/opre.26.2.305.

van Dijck J, Poell T, de Waal M (2018) *The Platform Society: Public Values in a Connective World.* Oxford University Press, Oxford.

Weber KM, Rohracher H (2012) Legitimizing research, technology and innovation policies for transformative change: combining insights from innovation systems and multi-level perspective in a comprehensive "failures" framework. *Research Policy* 41(6):1037–1047. doi:10.1016/j.respol.2011.10.015.

Williamson OE (1985) *The Economic Institutions of Capitalism: Firms, Markets, Relational Contracting.* Free Press, New York.