Latency and the Effective Capital Blueprint (ECB)

Charlie A. Johnson
Independent Research
research@charlieajohnson.com

January 31, 2026
Version 0.1

Abstract

Models of personal efficacy—from lifecycle consumption theory to portfolio choice—typically treat feasible action sets as expanding monotonically in wealth. Yet empirical patterns contradict this: entrepreneurs relocate and reverse course, high-net-worth individuals exhibit constrained portfolios, and operational scope narrows despite capital abundance. This paper introduces the *Effective Capital Blueprint* (ECB): the feasible action set an actor can execute given capital, coordination latency, and actor constraints. Our central result is that *coordination latency imposes non-compensable constraints on agency*: beyond latency thresholds, entire classes of coordination-critical actions are excluded from the feasible set regardless of capital availability. We show that coordination-critical actions requiring multiple feedback cycles exhibit threshold exclusion: small latency increases beyond action-specific thresholds cause discontinuous exit from the feasible set, with no capital-based restoration mechanism. In calibrated examples, tripling effective latency (from 2.2 to 7.7 days per cycle) reduces feasible iteration-intensive actions by approximately 65%, independent of capital levels. This challenges capital-centric models of personal efficacy and provides a formal foundation for observed clustering patterns that existing theories often treat as preference-driven or productivity-based. We derive testable predictions regarding portfolio composition, relocation reversals, and ecosystem persistence that distinguish latency-based from alternative explanations. Preliminary empirical tabulations on relocation events can be used to test these signatures at scale (see Section 9). These results imply that ecosystem persistence, geographic clustering, and relocation patterns reflect structural constraints on feasible action sets rather than amenity preferences or agglomeration externalities alone.

Notation

Symbol	Meaning
$\mathcal{X}$	Universe of possible actions
x	An action (or action class)
C	Available capital
$\mathbf{L}$	Latency vector $(L_1, \ldots, L_n)$
L_{eff}	Aggregate effective latency
$f(L_{\text{eff}})$	Increasing mapping from latency to cycle delay
$t_{\text{internal}}(x)$	Minimum internal processing time absent coordination delay
θ	Actor constraints (cognitive, temporal, expertise)
$\text{ECB}(C, \mathbf{L}, \theta)$	Effective Capital Blueprint (feasible action set)
C^*	Capital saturation threshold
L_x^*	Latency exclusion threshold for action x
$k(x)$	Feedback cycles required for action x
t_{cycle}	Cycle time including latency
δ_{PA}	Principal–agent coordination delay
δ_{coord}	Multi-party coordination overhead
$D(d)$	Context/domicile latency multiplier at context d
ε_x	Latency elasticity of feasibility of action x
T	Deadline window / time horizon
B	Actor cognitive bandwidth

1 Introduction

Why do actors with substantial capital exhibit sharply constrained agency? Conventional intuition treats capital as the primary constraint on action. Yet observed patterns suggest saturation: relocations reverse, portfolios narrow, and operational scope contracts—often without capital loss.

We propose that personal efficacy is bounded not by capital but by the *Effective Capital Blueprint* (ECB)—the feasible action set reachable within cognitive, temporal, and institutional limits. Coordination latency contracts this set nonlinearly and, for coordination-critical actions, *non-compensably.*

1.1 Contributions

1. **Theoretical:** Formalizes ECB as a set-valued function and establishes capital saturation and latency non-compensability under general conditions.
2. **Empirical:** Derives predictions, measurement proxies, and identification strategies distinguishing latency-based constraints from preference- or productivity-based explanations.
3. **Integrative:** Unifies transaction cost economics, institutional execution constraints, coordination theory, and spatial finance under a common feasibility-set lens.

1.2 Relation to Capital–Efficacy Literature

Prior work relates capital to efficacy via (i) direct affordability, (ii) signaling/access, and (iii) risk-bearing capacity. ECB framework does not contest these mechanisms. It identifies a fourth channel operating independently: *latency-constrained feasibility*. Under ECB, capital can be abundant and well-signaled yet still yield collapsed agency if coordination latency is high.

1.3 Paper Outline

Section 2 motivates the framework. Section 3 formalizes ECB and core definitions, then states assumptions. Section 4 structures coordination latency. Sections 5–7 develop the main results: capital saturation (Proposition 1), latency exclusion (Proposition 2), and non-compensability (Proposition 3), with comparative statics and control-theoretic grounding. Section 8 models context as a latency multiplier. Section 9 derives empirical predictions, robustness, and identification strategies. Sections 10–11 provide calibration/sensitivity and literature positioning. Sections 12–13 discuss extensions, welfare interpretation, and policy implications. Appendices include a referee-facing FAQ and implementation notes.

2 Conceptual Framework

2.1 Motivation

Two actors with identical capital can exhibit dramatically different agency across contexts. We argue this reflects latency-driven contraction of feasible action sets, not merely heterogeneous preferences or productivity.

2.2 Defining Personal Efficacy

Personal Efficacy is the capacity to convert intent into outcomes *quickly, predictably, and reversibly*, subject to real-world constraints. Efficacy derives from the *size and composition* of the feasible action set, not wealth in isolation.

3 The Effective Capital Blueprint (ECB)

3.1 ECB as a Set-Valued Function

Let $\mathcal{X}$ denote the universe of theoretically possible actions. Let $C \in \mathbb{R}_+$ be available capital, $\mathbf{L} \in \mathbb{R}^n_+$ a latency vector, and θ actor constraints. Define:

$$\mathrm{ECB}(C, \mathbf{L}, \theta) \subseteq \mathcal{X}$$

An action $x \in \mathcal{X}$ is in ECB iff it satisfies:

1. **Capital feasibility:** $\mathrm{cost}(x) \leq C$
2. **Latency feasibility:** $\mathrm{time}(x; \mathbf{L}) \leq T(\theta)$
3. **Cognitive feasibility:** $\mathrm{load}(x; \mathbf{L}) \leq B(\theta)$

ECB captures *reachability*, not theoretical affordability.

3.2 ECB Size Metric (Commitment)

For theory, define ECB size by cardinality:

$$\left|\mathrm{ECB}(C, \mathbf{L}, \theta)\right|$$

For empirics, use an outcome-weighted analogue:

$$V(\mathrm{ECB}) = \sum_{x \in \mathrm{ECB}} v(x)$$

Results extend to any monotonic aggregation over ECB.

3.3 Action Classification Preliminaries: Coordination-Critical Actions

An action x is **coordination-critical** if it requires $k(x) \geq k_{\min}$ decision–feedback cycles to complete within deadline T, where each cycle entails: (i) decision/action execution, (ii) observation of outcome/state, and (iii) adjustment based on feedback.

Define cycle time under latency $\mathbf{L}$:

$$t_{\text{cycle}}(x; \mathbf{L}) = t_{\text{internal}}(x) + f(L_{\text{eff}}),$$

where $f(\cdot)$ is increasing and $t_{\text{internal}}(x)$ is minimum internal processing time absent coordination delays.

An action is coordination-critical when:

$$k(x) \geq k_{\min} \quad \text{and} \quad k(x)\, t_{\text{cycle}}(x; \mathbf{L}) \approx T,$$

i.e., feasibility is tight with respect to coordination rounds rather than internal processing.

> Box 1: Concrete ECB Example
>
> Throughout this paper, we use archetypal contexts rather than specific geographic locations to emphasize that ECB mechanisms operate independently of particular places. The framework applies to any context pair with different latency multipliers $D(\cdot)$, whether driven by geography, institutional design, or organizational structure.
>
> Consider an early-stage founder with $C = \$2\text{M}$.
> **Abstract Hub (low-latency, high-density context; $D = 1.0$ by construction):** ECB includes daily synchronous iteration, rapid user testing ($\sim$3-day cycles), same-day legal iteration, and < 24h settlement responses. Illustratively $|\mathrm{ECB}| \approx 40$ concurrent action types.
> **Remote Periphery (high-latency, asynchronous context; $D \approx 3.5$):** Capital unchanged. ECB loses daily synchronous iteration (async substitutes fail), rapid testing cycles (7+ day minima), and fast legal feedback (multi-day response). Illustratively $|\mathrm{ECB}| \approx 22$ action types ($\sim$45% reduction).
> **Key insight:** nominal resources are unchanged, but entire action classes become unreachable due to latency alone.

3.4 Key Assumptions

A1 (Time–latency relationship): completion time is weakly increasing in effective latency:

$$\frac{\partial \mathrm{time}(x; \mathbf{L})}{\partial L_{\text{eff}}} \geq 0.$$

A2 (Load–latency relationship): monitoring/adjustment load increases with latency:

$$\frac{\partial \text{load}(x; \mathbf{L})}{\partial L_{\text{eff}}} \geq 0.$$

A3 (Finite actor constraints): $T(\theta) < \infty$ and $B(\theta) < \infty$.
A4 (Action discreteness): actions are discrete, distinct units (not infinitely divisible).
A5 (Delegation increases latency): deploying capital via delegation introduces coordination delays:

$$\delta_{\text{PA}} \geq 0, \quad \delta_{\text{coord}} \geq 0,$$

with at least one strict inequality for coordination-intensive actions.

4 Coordination Latency: Structure and Aggregation

4.1 Latency Vector

Coordination latency is modeled as $\mathbf{L} = (L_1, \ldots, L_n)$ representing communication delay, time-zone misalignment, institutional response, settlement, verification, etc.

4.2 Aggregation into Effective Latency

For tractability:

$$L_{\text{eff}} = \sum_{i=1}^{n} w_i L_i \quad \text{with } w_i > 0.$$

Alternative aggregations (Appendix) include multiplicative compounding $\prod_i (1 + \alpha_i L_i)$ and bottleneck $\max_i L_i$.

5 Main Results

5.1 Capital Saturation

Proposition 1 (Capital Saturation). *For fixed* $\mathbf{L}$ *and* θ*, there exists* $C^* < \infty$ *such that:*

$$\text{ECB}(C, \mathbf{L}, \theta) = \text{ECB}(C^*, \mathbf{L}, \theta) \quad \forall C > C^*.$$

Proof (deductive structure).

1. Let T denote the actor deadline window and B cognitive bandwidth (A3).

2. Define minimum completion time among latency-feasible actions:

$$t_{\min}(\mathbf{L}) = \min\{\text{time}(x; \mathbf{L}) : x \in \mathcal{X},\ \text{time}(x; \mathbf{L}) \leq T\}.$$

3. The maximum number of actions executable within T is bounded:

$$N_T \leq \left\lfloor \frac{T}{t_{\min}(\mathbf{L})} \right\rfloor.$$

4. Cognitive feasibility imposes a second bound. Let $b_{\min}$ be minimum attention per action:

$$N_B \leq \left\lfloor \frac{B}{b_{\min}} \right\rfloor.$$

5. Thus $|\mathrm{ECB}|$ is bounded above by $N_{\max} = \min(N_T, N_B)$, independent of C beyond affordability.
6. By A4 (action discreteness), the set of actions satisfying latency and cognitive constraints is finite; each has finite cost $\mathrm{cost}(x) < \infty$. Therefore:

$$C^* = \max\{\mathrm{cost}(x) : x \in \mathcal{X},\ \mathrm{time}(x; \mathbf{L}) \le T,\ \mathrm{load}(x; \mathbf{L}) \le B\}$$

   exists and is finite (maximum over finite set).
7. For any $C \ge C^*$, all latency- and cognitive-feasible actions are affordable, hence $\mathrm{ECB}(C, \mathbf{L}, \theta) = \mathrm{ECB}(C^*, \mathbf{L}, \theta)$. □

Key Takeaway

Capital saturates. Beyond C^*, more money does not expand what you can do—it only makes already-feasible actions easier to fund.

5.2 Action Type Preview (avoid forward reference)

We classify actions by latency sensitivity (formalized in Section 6):

- **Type I (latency-insensitive):** feasibility largely unaffected by L_{eff}.
- **Type II (latency-sensitive):** feasibility declines smoothly with L_{eff}, partially compensable by capital.
- **Type III (latency-critical):** discrete exclusion beyond threshold L_x^*, non-compensable by capital.

Proposition 3 establishes non-compensability for Type III actions.

5.3 Latency Exclusion

Proposition 2 (Latency Exclusion)**.** *Let $\mathbf{L}'$ differ from $\mathbf{L}$ such that $L'_{\text{eff}} \ge L_{\text{eff}}$. Then:*

$$\mathrm{ECB}(C, \mathbf{L}', \theta) \subseteq \mathrm{ECB}(C, \mathbf{L}, \theta).$$

Strict containment holds for non-trivial latency increases affecting at least one binding constraint for the given (C, θ).[1] *Moreover, there exist coordination-critical actions $X_{\text{crit}} \subset \mathcal{X}$ such that:*

$$X_{\text{crit}} \cap \mathrm{ECB}(C, \mathbf{L}', \theta) = \emptyset \quad \textit{even as } C \to \infty.$$

Proof (structure).

1. By A1, for any action x, $\mathrm{time}(x; \mathbf{L}') \ge \mathrm{time}(x; \mathbf{L})$ when $L'_{\text{eff}} \ge L_{\text{eff}}$.
2. If $x \in \mathrm{ECB}(C, \mathbf{L}, \theta)$, then $\mathrm{time}(x; \mathbf{L}) \le T$. Since $\mathrm{time}(x; \mathbf{L}') \ge \mathrm{time}(x; \mathbf{L})$, either:
   - $\mathrm{time}(x; \mathbf{L}') > T$, implying $x \notin \mathrm{ECB}(C, \mathbf{L}', \theta)$ (latency infeasibility), or
   - $\mathrm{time}(x; \mathbf{L}') \le T$ but $\mathrm{load}(x; \mathbf{L}') > B$ (by A2), implying $x \notin \mathrm{ECB}(C, \mathbf{L}', \theta)$ (cognitive infeasibility).

[1] Equality occurs only when the marginal latency increase affects solely non-binding components for the given (C, θ).

3. Therefore $\mathrm{ECB}(C, \mathbf{L}', \theta) \subseteq \mathrm{ECB}(C, \mathbf{L}, \theta)$.
4. For coordination-critical action classes X_{crit} requiring k cycles within T, if $k(t_{\mathrm{base}} + f(L'_{\mathrm{eff}})) > T$ but $k(t_{\mathrm{base}} + f(L_{\mathrm{eff}})) \leq T$, then $X_{\mathrm{crit}} \cap \mathrm{ECB}(C, \mathbf{L}', \theta) = \emptyset$ regardless of C (time constraint binds independently of capital). □

5.4 Non-Compensability

Proposition 3 (Non-Compensability for Coordination-Critical Actions)**.** *Let X_{crit} be actions requiring k decision–feedback cycles within deadline T. Define:*

$$t_{\mathrm{cycle}}(L_{\mathrm{eff}}) = t_{\mathrm{base}} + f(L_{\mathrm{eff}}),$$

where $f(\cdot)$ is strictly increasing (A1–A2). Feasibility requires:

$$k\, t_{\mathrm{cycle}}(L_{\mathrm{eff}}) \leq T.$$

Define threshold L^* by:

$$k\,(t_{\mathrm{base}} + f(L^*)) = T.$$

Then for $L_{\mathrm{eff}} > L^*$:

$$X_{\mathrm{crit}} \cap \mathrm{ECB}(C, \mathbf{L}, \theta) = \emptyset \quad \forall C < \infty.$$

Proof (deductive structure).

1. If $L_{\mathrm{eff}} > L^*$, then $k(t_{\mathrm{base}} + f(L_{\mathrm{eff}})) > T$, so direct execution violates the deadline.
2. **Capital deployment via delegation:** the only plausible capital-based restoration mechanism is offloading work to hired agents (delegation/subcontracting).
3. By A5, delegation introduces nonnegative delays: principal–agent delay $\delta_{\mathrm{PA}}(k) \geq 0$ and coordination overhead $\delta_{\mathrm{coord}}(k) \geq 0$. For coordination-intensive actions (high k), at least one is strictly positive.
4. Effective cycle time under delegation is:

$$t'_{\mathrm{cycle}} = t_{\mathrm{base}} + f(L_{\mathrm{eff}}) + \delta_{\mathrm{PA}}(k) + \delta_{\mathrm{coord}}(k).$$

5. Since $\delta_{\mathrm{PA}} + \delta_{\mathrm{coord}} > 0$ for Type III actions:

$$k\, t'_{\mathrm{cycle}} > k(t_{\mathrm{base}} + f(L_{\mathrm{eff}})) > T,$$

   so delegation worsens feasibility.
6. Thus no finite capital can restore feasibility once $L_{\mathrm{eff}} > L^*$. □

Key Takeaway

For coordination-critical actions, latency thresholds create hard walls. Once crossed, no amount of capital restores feasibility.

5.5 Comparative Statics Summary

Parameter	Effect on ECB size	Magnitude	Most affected
$\uparrow C$	$+$ below C^*, 0 above C^*	Saturating	Type I initially
$\uparrow L_{\text{eff}}$	$-$ monotone	Nonlinear / threshold	Type III > II > I
$\uparrow B$	$+$	Linear–sublinear	Type II
$\uparrow T$	$+$	Linear	Type III
$\uparrow D(d)$	$-$ via latency	Multiplicative	All

6 Action Taxonomy and Measurement

6.1 Latency Elasticity

Let $P_x = \Pr[x \in \text{ECB}(C, \mathbf{L}, \theta)]$ and define:

$$\varepsilon_x = -\left(\frac{dP_x}{dL_{\text{eff}}}\right)\left(\frac{L_{\text{eff}}}{P_x}\right).$$

6.2 Formal Classification

Define latency threshold:

$$L_x^* = \inf\{L_{\text{eff}} : x \notin \text{ECB}(C, \mathbf{L}, \theta) \text{ for } C \text{ sufficiently large}\}.$$

Type I: $L_x^* = \infty$ and $\varepsilon_x \approx 0$. Type II: $0 < L_x^* < \infty$ and ε_x finite. Type III: $L_x^* < \infty$ and $\lim_{L_{\text{eff}} \to L_x^*} \varepsilon_x = \infty$.

6.3 Empirical Operationalization

Estimate latency sensitivity using completion probabilities:

$$\hat{\varepsilon}_x = -\frac{\Delta \Pr[x \text{ completed}]}{\Delta L_{\text{eff}}} \cdot \frac{\overline{L_{\text{eff}}}}{\overline{\Pr[x \text{ completed}]}}.$$

Concrete application (worked example). Consider a "contract negotiation" action requiring $k = 5$ rounds of revision within a $T = 30$ day window.

- Low-L regime ($D = 1.0$): $L_{\text{eff}} = 2.2$ days $\Rightarrow$ completion time ≈ 11 days; completed ≈ 0.95 within T.
- High-L regime ($D = 3.5$): $L_{\text{eff}} = 7.7$ days $\Rightarrow$ completion time ≈ 38.5 days; completed ≈ 0.05 within T.

Estimated elasticity:

$$\hat{\varepsilon}_x = -\left[\frac{0.05 - 0.95}{7.7 - 2.2}\right] \cdot \frac{(7.7 + 2.2)/2}{(0.95 + 0.05)/2} = -\left[\frac{-0.90}{5.5}\right] \cdot \left[\frac{4.95}{0.5}\right] = 0.164 \cdot 9.9 \approx 1.6.$$

This places the action in Type II (finite positive elasticity between 0.1 and 5). For actions requiring $k \geq 10$ rapid iterations (Type III), completion probabilities fall discontinuously to near-zero, implying $\hat{\varepsilon}_x \to \infty$.

6.4 Algorithmic Type Classification (implementable)

Algorithm 1: Action Type Classification

Input: action x, panel data of completion outcomes across observed L_{eff} values
Output: Type I / II / III classification

1. Estimate $P(x \mid L_{\text{eff}})$ using logistic regression or a flexible spline model.
2. Compute numerical derivative dP/dL_{eff} at median L_{eff}.
3. Compute $\hat{\varepsilon}_x$ at median L_{eff}.
4. If $\hat{\varepsilon}_x < 0.1$, classify Type I.
5. Else if decline is smooth with finite $\hat{\varepsilon}_x$, classify Type II.
6. Else if a discontinuity is detected (e.g., likelihood ratio test comparing threshold vs smooth models), classify Type III.

Computational complexity: For N observations and G grid points in L_{eff}, step 1 requires $O(N \log N)$ for logistic regression (typical solvers) or $O(NG)$ for spline fitting. Steps 2–6 are $O(G)$. Total runtime is $O(N \log N + NG)$, feasible for typical empirical datasets.

7 Formal Foundations: Control-Theoretic Grounding

ECB connects formally to control theory via stability and gain–bandwidth limitations.

Model efficacy as a feedback control system with observation delay $\tau = L_{\text{eff}}$ (feedback delay in the control-theoretic sense). A standard inverse-delay gain limitation (see, e.g., Åström & Murray; Franklin, Powell & Emami-Naeini) implies that for frequency ω when $\omega\tau > 1$:

$$|G_{\max}(\omega)| \leq \frac{1}{\omega\tau}.$$

Type III actions have high natural frequency ω_n (tight control requirements). When $L_{\text{eff}} \cdot \omega_n > 1$, the system becomes effectively unstable/uncontrollable regardless of control effort (capital), providing a control-theoretic foundation for Proposition 3.

8 Context as a Latency Multiplier

Let $d \in \Omega$ denote a context. Let $D : \Omega \to \mathbb{R}_+$ map context to a systematic latency multiplier:

$$L_i(d) = D(d) \cdot L_i^{\text{base}} + \varepsilon_i(d),$$

with aggregate:

$$L_{\text{eff}}(d) = D(d) \cdot L_{\text{eff}}^{\text{base}} + \sum_i w_i \varepsilon_i(d).$$

Thus $\text{ECB}(C, \mathbf{L}(d_1), \theta)$ can differ dramatically from $\text{ECB}(C, \mathbf{L}(d_2), \theta)$ even when C and θ are identical.

9 Empirical Implications, Identification, and Robustness

9.1 Predictions

(i) Portfolio narrowing with higher L_{eff} even at high C; (ii) turnover suppression; (iii) threshold exclusion for Type III near L_x^*; (iv) reversibility premia rise with L_{eff}; (v) relocation reversals when L_{eff} increases.

9.2 Identification

Within-actor relocation event study / DID:

$$\Delta\text{ECB} = \beta_0 + \beta_1 \Delta L_{\text{eff}} + \beta_2 \Delta C + \text{controls} + \epsilon,$$

expecting $\beta_1 < 0$.

9.3 Robustness: Addressing Alternative Explanations

- **Reverse causality:** declining ECB might precede relocation. Test pre-trends; ECB predicts flat pre-trends and sharp post-change. Placebo: relocations between similar-L_{eff} contexts should show no ECB change.
- **Omitted variables:** unobserved shocks could correlate with relocation and ECB. Restrict to plausibly forced relocations. Use institutional shocks as instruments for L_{eff}.
- **Measurement error:** self-reports noisy. Validate with objective proxies (transaction logs; calendar density; contract cycle times).

9.4 Power (Illustrative)

To detect $\beta_1 = -0.20$ with power 0.80, $\alpha = 0.05$, within-actor correlation $\rho = 0.6$, residual variance $\sigma^2 = 0.04$ requires approximately $N \approx 85$ relocation events. Threshold-focused Type III tests may require $N \approx 40$ under sharp discontinuity.

10 Calibration and Sensitivity

Tripling L_{eff} from 2.2 to 7.7 days per cycle reduces feasible iteration capacity from ~ 13 to ~ 4 cycles per month. Type III actions requiring more than 4 iterations become infeasible independent of C, consistent with the $\sim 65\%$ reduction anchor used in the abstract. Sensitivity: L^* scales linearly in T and t_{base}, and inversely in k.

11 Relation to Existing Literature and Gap

ECB bridges transaction cost economics (Coase; Williamson), institutional economics (North; Acemoglu & Johnson), coordination theory (Thompson; Galbraith; Malone & Crowston), geography of finance (Coval & Moskowitz; Petersen & Rajan), remote work (Bloom et al.; Emanuel & Harrington), and agglomeration (Duranton & Puga; Ellison et al.).

Gap this paper fills (specific)

Existing frameworks face three limitations ECB resolves:

1. **Transaction cost theory** treats costs as compensable with resources; ECB shows latency creates hard constraints where delegation increases the constraint via δ_{PA} and δ_{coord}.
2. **Geography of finance** attributes distance effects to information asymmetry; ECB shows direct feasibility loss even with perfect information.
3. **Agglomeration theory** explains clustering via productivity spillovers; ECB adds a non-additive mechanism: departure is infeasible for Type III actors when alternatives have $L_{\text{eff}} > L^*$.

Additional distinction (remote work boundary): Remote work evidence can show output-per-hour remains high for some tasks, yet ECB predicts a bifurcation: Type I tasks remain feasible, while Type III tasks can become infeasible (discrete exclusion) even absent a productivity decline.

12 Extensions

(i) Obligation load Ω in $\text{ECB}(C, \mathbf{L}, \theta, \Omega)$; (ii) dynamic θ and endogenous network formation reducing some L_i; (iii) multi-actor ECB with internal latency.

12.1 Actor heterogeneity in ECB sensitivity

Actor class	Dominant action types	ECB sensitivity to L_{eff}	Typical C^*
Passive investors	Type I	Low	Low
Content creators	Type I	Low	Medium
Service professionals	Type II	Medium	Medium
Entrepreneurs	Type III	High	High
Active portfolio managers	Type II–III	High	Very high

Prediction: relocation effects concentrate among high-k actors with near-zero effects for low-k populations.

13 Conclusion

Personal efficacy is constrained by ECB—the feasible action set reachable within cognitive, temporal, and institutional limits. Capital saturates in its ability to expand ECB. Coordination latency contracts ECB and, for coordination-critical actions, does so non-compensably: once latency crosses thresholds, no finite capital restores feasibility.

13.1 Welfare and Policy Implications

Connection to utility theory: represent preferences as $U(c, |\text{ECB}|)$ with $\partial U/\partial c > 0$ and $\partial U/\partial|\text{ECB}| > 0$.[2] Even if consumption c remains constant, welfare decreases when $|\text{ECB}|$ contracts.

[2] We assume separability for expositional clarity. More generally, $U(c, \text{ECB})$ could exhibit complementarity—consumption value may depend on feasible actions. This would strengthen rather than weaken the welfare-loss logic under ECB contraction.

Define willingness-to-pay $v(x)$ to restore action x to feasibility, and welfare loss from a latency increase:

$$\Delta W = \sum_{x \in \mathrm{ECB}(\mathbf{L}_{\mathrm{low}}) \setminus \mathrm{ECB}(\mathbf{L}_{\mathrm{high}})} v(x).$$

Policy: reducing L_{eff} (institutional responsiveness, settlement speed, permitting/court latency) may expand feasible action sets more than capital injections alone; remote-work policies may work for Type I activities but fail for Type III.

References

[1] Acemoglu, D., & Johnson, S. (2005). Unbundling institutions. *Journal of Political Economy*, 113(5), 949–995.

[2] Åström, K. J., & Murray, R. M. (2008). *Feedback Systems: An Introduction for Scientists and Engineers*. Princeton University Press.

[3] Bloom, N., Liang, J., Roberts, J., & Ying, Z. J. (2015). Does working from home work? Evidence from a Chinese experiment. *Quarterly Journal of Economics*, 130(1), 165–218.

[4] Coase, R. H. (1937). The nature of the firm. *Economica*, 4(16), 386–405.

[5] Coval, J. D., & Moskowitz, T. J. (1999). Home bias at home: Local equity preference in domestic portfolios. *Journal of Finance*, 54(6), 2045–2073.

[6] Coval, J. D., & Moskowitz, T. J. (2001). The geography of investment: Informed trading and asset prices. *Journal of Political Economy*, 109(4), 811–841.

[7] Duranton, G., & Puga, D. (2004). Micro-foundations of urban agglomeration economies. In *Handbook of Regional and Urban Economics* (Vol. 4). Elsevier.

[8] Ellison, G., Glaeser, E. L., & Kerr, W. R. (2010). What causes industry agglomeration? Evidence from coagglomeration patterns. *American Economic Review*, 100(3), 1195–1213.

[9] Emanuel, N., & Harrington, E. (2023). Working remotely? Selection, treatment, and the market for remote work. *Journal of Labor Economics* (forthcoming).

[10] Franklin, G. F., Powell, J. D., & Emami-Naeini, A. (2015). *Feedback Control of Dynamic Systems* (7th ed.). Pearson.

[11] Galbraith, J. R. (1973). *Designing Complex Organizations*. Addison-Wesley.

[12] Malone, T. W., & Crowston, K. (1994). The interdisciplinary study of coordination. *ACM Computing Surveys*, 26(1), 87–119.

[13] Milgrom, P., & Roberts, J. (1992). *Economics, Organization and Management*. Prentice Hall.

[14] North, D. C. (1990). *Institutions, Institutional Change and Economic Performance*. Cambridge University Press.

[15] Petersen, M. A., & Rajan, R. G. (2002). Does distance still matter? The information revolution in small business lending. *Journal of Finance*, 57(6), 2533–2570.

[16] Spence, M. (1973). Job market signaling. *Quarterly Journal of Economics*, 87(3), 355–374.

[17] Thompson, J. D. (1967). *Organizations in Action*. McGraw-Hill.

[18] Williamson, O. E. (1985). *The Economic Institutions of Capitalism*. Free Press.

A Appendix D: Addressing Common Questions

Q1: Don't preferences explain relocation reversals?

Preference-based explanations predict correlation between stated preferences and reversals. ECB predicts reversals even when preferences are stable, distinguished by Type III exclusion patterns.

Q2: Can't technology (video calls, async tools) offset latency?

For Type I/II actions, partially. For Type III actions requiring tight control loops, fundamental control-theoretic limits bind (Section 7), producing testable heterogeneity.

Q3: Why doesn't capital enable technology adoption to reduce L_{eff}?

Technology reduces some components (communication delay) but not others (timezone misalignment, institutional response, settlement). For coordination-critical actions, remaining components can still push L_{eff} beyond L^*.

Q4: Is this just "you can't be in two places at once"?

No. The contribution is quantitative: how much latency before which actions become infeasible, and why capital cannot compensate. Non-compensability (Proposition 3) contradicts the intuition that "money solves problems."

Q5: Does this apply to firms or only individuals?

It generalizes. Firms face internal coordination latency; large organizations may have higher effective L_{eff} than small teams, yielding coordination diseconomies for iteration-intensive work.