The Reach of Monetary Policy

Evangelos D. Stravelas*

September 2026

Conference submission version. An updated version will follow shortly.

Abstract

No monetary authority operating across heterogeneous regions can simultaneously coordinate inflation expectations, transmit policy to the real economy, and stabilize inflation uniformly. Three logically independent failures build the barrier: a scalar interest rate cannot span a multi-dimensional structural state space; bounded information capacity forces uneven credibility judgments across regions; and regional welfare losses are not recovered by aggregate gains elsewhere. Each failure rests on a single physical primitive—finite information bandwidth, irreducible regional heterogeneity, and strictly concave regional welfare—and each survives when the defining conditions of the other two are absent. Under the canonical benchmark, the impossibility takes closed form as two thresholds: an informational ceiling above which transmission collapses, and a coordination threshold above which expectation divergence becomes self-reinforcing. Under rational expectations, the coordination threshold collapses to zero. No institutional architecture expands the feasibility set.

JEL codes: E52, E58, D83, D84, F41.

Keywords: Monetary policy; monetary reach; expectation coordination; policy transmission; regional heterogeneity; bounded attention; rational inattention; optimum currency areas.

*e.stravelas@gmail.com. The author thanks seminar participants for comments. All errors are the author's own. Data availability: No data were used in this theoretical study.

1. Introduction

No monetary authority operating in a heterogeneous jurisdiction can simultaneously coordinate inflation expectations, transmit policy to the real economy, and stabilize outcomes uniformly across regions. The central bank commands a scalar interest rate against $K \geq 2$ dimensions of structural variation. This dimensional gap is the operative constraint: any spatial divergence in regional natural rates—measured by their cross-sectional standard deviation $F \equiv \sqrt{\mathrm{Var}_i(r_i^*)} \geq 0$—makes joint achievement mathematically infeasible. The impossibility is a regime viability result, not a trade-off; at exact tolerance zero, the feasibility set is strictly empty for every $F > 0$. The impossibility manifests across three policy objectives formally defined as tolerance bounds on equilibrium outcomes: *Consistent Expectation Coordination* CE(ε) requires $\mathrm{Var}_i(\mathbb{E}_i[\pi_{t+1}]) \leq \varepsilon$ (belief alignment, not strategic complementarity); *Effective Policy Transmission* PE(δ) requires transmission span $\kappa(F)\sigma(F) \geq \delta$; and *Uniform Regional Stabilization* US(η) requires $\mathrm{Var}_i(\pi_i - \pi^*) \leq \eta$.

Three structural primitives suffice for the universal impossibility: bounded information capacity (P1) and irreducible natural-rate heterogeneity (P2) enforce the absolute physical transmission and coordination limits, while strictly concave regional welfare (P3) isolates the resulting unrecoverable aggregate welfare deficit. First, the scalar projection embeds a non-trivial kernel by the rank-nullity theorem; unspanned variance severs expectation coordination from policy transmission (CE–PE). Second, irreducible heterogeneity floors spatial inflation variance at strictly positive limits, fracturing transmission from uniform stabilization (PE–US). Third, strictly concave regional welfare converts residual belief dispersion into an absolute Jensen welfare deficit; losses in deprived regions are not recovered by gains elsewhere, permanently separating coordination and stabilization (CE–US). Three bilateral fractures unify

into a single dimensional barrier (Theorem 4.1). The barrier is geometric, not frictional: it follows from the rank-deficiency of the scalar policy projection and the unspanned spatial variance that deficiency forces, independent of the specific source of short-run non-neutrality.

The Tinbergen instrument-counting principle is necessary but not sufficient: the binding constraint is the dimensional rank of the scalar projection against the $K \geq 2$ structural state space, not the count of policy objectives. The canonical New Keynesian optimal commitment rule (Woodford, 2003) solves for a representative inflation path that exists only at $F = 0$; at any $F > 0$, spatial heterogeneity is a dimensional constraint, not a perturbation of the homogeneous benchmark. The classical Mundell-Fleming trilemma is the degenerate first-order linearization of this impossibility at the homogeneous limit (Theorem 4.2); optimum currency area theory characterizes the desirability of monetary union formation, not the dimensional constraints binding the authority once formed (Mundell, 1961; McKinnon, 1963; Kenen, 1969). Regional transmission elasticities vary with trade exposure, holding natural rates fixed (Hottman and Reyes-Heroles, 2024). Enforcing uniform transmission isolates the impossibility strictly to natural-rate heterogeneity and informational rank-deficiency: centralized coordination fails on this basis alone, requiring no asymmetry in transmission. Dimensional rank-deficiency severs the public-signal precision lever before it reaches agents and prevents the public signal from spanning the spatial state: Morris and Shin (2002)'s strategic coordination mechanism is structurally inoperative and Angeletos and Lian (2018)'s common-knowledge channel is geometrically bypassed; the CE failure is physical, not strategic. The rational inattention literature (Sims, 2003; Mackowiak and Wiederholt, 2009) treats bounded capacity as an optimized choice over a fixed signal space; bounded capacity here is a physical primitive (P1) that binds

even the rationally inattentive optimum, placing the CE failure strictly outside capacity-misallocation frameworks. The PE failure revises Barro and Gordon (1983): dimensional rank-deficiency intrinsically generates unequal credibility across regions, rendering the credibility problem spatial as well as temporal (Clarida, Galí, and Gertler, 1999). The impossibility defines the cross-regional feasibility set within which heterogeneous-agent redistribution and liquidity channels operate (Cúrdia and Woodford, 2010; Auclert, 2019; Kaplan, Moll, and Violante, 2018); within-region agent heterogeneity and the spatial impossibility are orthogonal dimensions. Any multi-instrument design must operate within the geometric feasibility limits established here; the scalar-instrument restriction is the unique binding constraint (Alexandrov and Brunnermeier, 2026).

The impossibility partitions the geometric space across two macroscopic boundaries. The aggregate transmission ceiling F_1 isolates the threshold above which unspanned spatial variance forces aggregate attention to release the capacity constraint, degrading policy transmission asymptotically to zero; base transmission elasticities (κ_0, σ_0) are point-identified from aggregate data and spatial extraction capacity $(h_a, \bar{H})$ delivers set identification bounds. The coordination threshold $F_3^* = \inf\{F \in [0, F_1) : \mu_c(F) \geq 1\}$ is the first crossing at which the expectation-divergence coefficient $\mu_c(F) = a_\pi^2 \bar{X}(F)^2$ reaches unity; above it expectation dispersion is geometrically expansive and CE(ε) fails for every tolerance; the attenuation rate C'_λ that pins F_3^* is set-identified from cross-sectional regional forecast errors. Under exact rational expectations the feedback coefficient is $\mu_c^{\mathrm{RE}} = a_\pi^2$, independent of F; no interior fragmentation threshold exists, and whether dispersion converges or diverges is decided by $(\kappa_0, \sigma_0, \beta)$ alone.

For any $F > 0$, dimensional rank constraints enforce strict lower bounds $A(F, t) > 0$ and $\mathrm{Var}_i(\pi_i - \pi^*) > 0$, algebraically eliminating exact simultaneous fulfillment at $\varepsilon =$

$\eta = 0$. Above F_3^*, the expectation-divergence coefficient satisfies $\mu_c(F) > 1$; CE(ε) fails for every positive tolerance $\varepsilon > 0$, making the impossibility tolerance-free on this interval. The central bank operates within geometric limits that no institutional architecture can expand.

2. Economic Environment and Structural Conditions

2.1 The Economy

The heterogeneous jurisdiction distributes continuously across the spatial measure $i \in [0, 1]$; each region maps to a structural state vector $\theta_i \in \Theta \subset \mathbb{R}^K$. $\mathrm{Var}_i(\cdot)$ denotes the cross-sectional variance over the spatial measure at a fixed period t; $\mathrm{Var}(\cdot)$ denotes the intertemporal variance of the stochastic processes. The cross-sectional standard deviation of regional natural rates,

$$F \equiv \sqrt{\mathrm{Var}_i(r_i^*)} \geq 0, \tag{1}$$

enters the impossibility through the dimensionless term $\gamma_u^2 F^2 = \gamma_u^2 \mathrm{Var}_i(r_i^*)$, where $\gamma_u > 0$ is any nonzero transmission scalar carrying units of output gap per unit natural rate. The symbol γ_u is deliberately generic: it is not γ_A, γ_I, γ_c, or a_π, and no universal result below uses its value. The three immutable physical primitives govern the structural geometry of the jurisdiction independently of any policy protocol or expectation-formation mechanism: *Bounded Capacity* (P1): aggregate spatial filtration capacity is finite, $\bar{H} < \infty$; *Irreducible Heterogeneity* (P2): $\mathrm{Var}_i(r_i^*) > 0$, independently of observation technology or capacity (Friedman, 1968); *Concave Welfare* (P3): regional utility satisfies $u''(\bar{c}, \bar{n}) < 0$. The monetary instrument is restricted to a scalar $R_t \in \mathbb{R}$ at each date t; infinite-dimensional fiscal transfer

vectors exit the admissible domain $\mathfrak{M}(\mathfrak{D})$ by violating Condition (C) of Definition 2.2. Rejecting P1 formally asserts infinite Shannon channel bandwidth (Cover and Thomas, 1991; Sims, 2003), a configuration outside $\mathfrak{D}$.

Assumption 2.1 (Filtration Separation and Initialization — Domain Restriction)**.** The spatial cross-section initializes at its invariant measure at $t = 0$. The idiosyncratic filtration $\{\Omega_{it}^{\text{idio}}\}_{i\in[0,1]}$ operates independent of the aggregate probability space $(\Omega, \mathcal{F}, \mathbb{P})$. For strictly exogenous, mean-zero idiosyncratic sequences v_{it} drawn independently across regions, the Exact Law of Large Numbers (Judd, 1985) enforces $\int_0^1 v_{it}\, di = 0$ almost surely at every finite t.

Assumption 2.2 (Central Bank Attention — Domain Restriction)**.** The central bank filtration operator satisfies $\bar{H}_{CB} < \infty$.

Assumption 2.3 (Drifting Central Bank Type (A-RW) — Benchmark)**.** The central bank's unobserved type evolves as $\theta_{CB,t+1} = \theta_{CB,t} + \nu_t$, with $\nu_t \sim \mathcal{N}(0, \sigma_\nu^2)$ and $\sigma_\nu^2 > 0$.

Assumption 2.4 (Positive-Measure Optimizing Subclass — Domain Restriction)**.** The subset of the continuum resolving the interior extraction problem constitutes a strictly positive Lebesgue measure $m \in (0, 1]$.

Assumption 2.5 (Monetary Transmission Primitives — Domain Restriction)**.** MTP-1 (*Price response to real activity*): $\kappa(F) > 0$ links inflation deviations to output-gap deviations. MTP-2 (*Output response to real financing cost*): $\sigma(F) > 0$ links output-gap deviations to the unspanned real financing cost gap. MTP-3 (*Price predetermination*): π_{it} is measurable with respect to $\Omega_{i,t-1}$. The spatial mappings accommodate Calvo (1983) pricing, Taylor contracts, IS-LM, and post-Keynesian markup dynamics. Corollary 3.2 already extends the transmission channel to any continuous $\kappa(F), \sigma(F)$ with $\kappa(F)\sigma(F) > 0$, not to the Calvo parametrization

specifically; the dispersion results of Section 3 inherit the same genericity, since their derivation uses nothing about κ_0, σ_0 beyond positivity. What Section 3 requires is not a particular pricing friction but a particular expectations protocol: the recursive-anchoring rule of Assumption 3.10. Under full rational expectations, equilibrium selection forces the coefficient on any unstable forecast path to zero; the explosive channel does not survive that limit.

Assumption 2.6 (Information-Transmission Scaling — Benchmark)**.** The regional transmission coefficients scale proportionally to aggregate spatial attention: $\kappa_i(F) = \kappa_0 \cdot h_{\text{agg}}(F)/h_a$ and $\sigma_i(F) = \sigma_0 \cdot h_{\text{agg}}(F)/h_a$. The aggregate transmission scalars are $\kappa(F) \equiv \int_0^1 \kappa_i(F)\, di$ and $\sigma(F) \equiv \int_0^1 \sigma_i(F)\, di$. The proportional scaling belongs to the generic class satisfying $\kappa(F)\sigma(F) > 0$ for all finite F; it provides the closed-form boundary evaluation within this class.

The log-linearized equations under Calvo pricing instantiate MTP-1 through MTP-3:

$$\pi_{it} = \beta \mathbb{E}_{i,t-1}[\pi_{it}] + \kappa_i(F)\mathbb{E}_{i,t-1}[y_{it}^{\text{gap}}], \tag{2}$$

$$y_{it}^{\text{gap}} = \mathbb{E}_{it}[y_{i,t+1}^{\text{gap}}] - \sigma_i(F)\big(R_t - \mathbb{E}_{it}[\pi_{i,t+1}] - r_i^*\big). \tag{3}$$

The natural rate is $r_i^* = \rho + \psi\theta_i$, with $\text{Var}_i(r_i^* - r^*) = \psi^2 \,\text{Var}_i(\theta_i) > 0$ for any spatial variation $F > 0$.

The monetary authority minimizes the cross-sectional aggregate loss functional, which penalizes deviations of regional inflation from the target π^* and regional output gaps from zero:

$$L_t = \int_0^1 \Big[(\pi_{it} - \pi^*)^2 + \alpha\,(y_{it}^{\text{gap}})^2\Big]\, di, \tag{4}$$

where $\alpha > 0$ is the relative weight on output stabilization and the target $\pi^* \geq 0$ is

the exogenous scalar defined in condition (D6).

Definition 2.1 (Admissible Equilibrium)**.** An admissible equilibrium within $\mathfrak{D}$ constitutes three geometric requirements. (A) *Physical Capacity Bound*: The aggregate attention integral satisfies $\int_0^1 h_i \, di \leq h_a < \infty$. (B) *Measurability and Optimality*: The scalar aggregate policy function $R_t = \varphi(\Omega_{CB,t})$ uniquely minimizes (4); every localized expectation field ξ_{it} is strictly measurable and adapted to the bounded individual filtration Ω_{it}. (C) *Structural Consistency*: Equilibrium objects iterate over the exogenous environment, governed by the structural primitives independently of the expectation-formation protocol.

Lemma 2.1 (Existence of Admissible Equilibrium)**.** *An admissible equilibrium in* $\mathfrak{M}(\mathfrak{D})$ *exists.*

Proof. The feasible attention set $\mathcal{H} \equiv \{h \in L^\infty([0,1]) : 0 \leq h(i) \leq h_a \ \mu\text{-a.e.}\}$ is norm-bounded, weak*-closed, and convex; Banach–Alaoglu delivers weak*-compactness. The loss (4) is strictly convex in R_t with $L_t \to \infty$ as $|R_t| \to \infty$; Berge's Maximum Theorem delivers a continuous, single-valued optimal policy correspondence $\Phi : \mathcal{H} \to \mathcal{R}$ on the compact interval $\mathcal{R} \subset \mathbb{R}$. The optimal attention correspondence $\Psi : \mathcal{R} \rightrightarrows \mathcal{H}$ is non-empty-valued, convex-valued, and weak*-upper hemicontinuous: metrizable by separability of $L^1([0,1])$, sequential upper hemicontinuity follows from dominated convergence preserving the capacity constraint at any weak*-convergent limit. Kakutani–Glicksberg–Fan delivers a fixed point $(R^*, h^*) \in \mathcal{R} \times \mathcal{H}$, securing Conditions (A)–(C) of Definition 2.1. The infinite-horizon value function exists by the Theorem of the Maximum applied recursively (Stokey, Lucas, and Prescott, 1989). ■

Lemma 2.2 (Aggregate Output Gap: The Projection Identity)**.** *Under Condition (B) of Definition 2.1, MTP-2, MTP-3, Assumption 2.6, Assumption 2.2, and* $\alpha > 0$*, the*

cross-sectional mean output gap equals the central bank's aggregate forecast error:

$$\bar{y}^{\text{gap}}(t) \equiv \int_0^1 y_{it}^{\text{gap}}\, di = \bar{b}_t - \mathbb{E}\big[\bar{b}_t \mid \Omega_{CB,t}\big],$$
$$\bar{b}_t \equiv \int_0^1 \Big(\mathbb{E}_{it}[y_{i,t+1}^{\text{gap}}] + \sigma(F)\mathbb{E}_{it}[\pi_{i,t+1}] + \sigma(F) r_i^*\Big)\, di. \tag{5}$$

The mean gap vanishes if and only if the aggregate IS position is measurable with respect to the bank's own filtration. Bounded central-bank capacity (P1, Assumption 2.2) denies that measurability in general, so no zero-average result follows from optimality alone.

Proof. By MTP-3, π_{it} is $\Omega_{i,t-1}$-measurable, so $\int_0^1 (\pi_{it} - \pi^*)^2\, di$ in (4) is fixed at the date R_t is chosen and drops out of the first-order condition for R_t. By Assumption 2.6, $\sigma_i(F) = \sigma(F)$ is common across i; (3) gives $y_{it}^{\text{gap}} = b_i - \sigma(F)R_t$ with b_i the integrand of $\bar{b}_t$. Condition (B) requires $R_t = \varphi(\Omega_{CB,t})$, so the bank minimizes $\alpha\,\mathbb{E}\big[\int_0^1 (b_i - \sigma(F)R_t)^2 di \mid \Omega_{CB,t}\big]$ over $\Omega_{CB,t}$-measurable R_t; with $\alpha > 0$ and $\sigma(F) > 0$ the objective is strictly convex, and the unique minimizer is the conditional projection $R_t^* = \mathbb{E}[\bar{b}_t \mid \Omega_{CB,t}]/\sigma(F)$. Substituting gives (**??**). ■

Corollary 2.1 (Zero Average Gap and Zero Average Inflation — Benchmark)**.** *Impose* aggregate policy sufficiency*: $\bar{b}_t$ is $\Omega_{CB,t}$-measurable. Then $\bar{y}^{\text{gap}}(t) = 0$ at every date; if in addition Assumption 3.9 holds, averaging* (2) *under MTP-3 gives $(1-\beta)\bar{\pi}(t) = 0$, hence $\bar{\pi}(t) = 0$. The restriction is informational, not optimizing, and it is not free: Theorem 4.2 exhibits an admissible homogeneous environment in which $\sigma_e^2 > 0$ makes it fail and the realized aggregate gap $\sigma_0[(1-\phi)\zeta_t - \phi e_t]$ is strictly random. Every result below that uses a zero aggregate gap cites this corollary, never Lemma 2.2.*

2.2 Structural Conditions and Admissible Class

Definition 2.2 (Admissible Economic Domain)**.** Analysis is restricted to the canonical domain of scalar-instrument monetary architectures. Let the Admissible Domain $\mathfrak{D}$ be defined by the intersection of the foundational primitives (P1–P3) and the strict dimensional constraints (D1–D3, D5–D7). Itemized: **(D1)** the structural state space has dimension $K \geq 2$, $\theta_i \in \Theta \subset \mathbb{R}^K$ for every region i; **(D2)** the monetary instrument is restricted to a scalar $R_t : \mathbb{R}^K \to \mathbb{R}$, for any measurable map, linear or non-linear; **(D3)** consequently, every admissible scalar projection δ_t leaves a residual of dimension $K - 1 \geq 1$ and admits no right inverse; Corollary 4.1 proves the exact residual this forces on the policy-relevant target directly, by Hilbert-space projection onto the one-dimensional subspace spanned by the scalar instrument, without appeal to rank-nullity for a general nonlinear map. **(D5)** the idiosyncratic innovation ε_{it} is i.i.d. across (i, t) and independent of $\Omega_{i,t-1}$, drawn from a fixed distribution such that the resulting cross-sectional residual belongs to a location-scale family, and it enters realized regional outcomes additively, $\pi_{it} = \tilde{\pi}_{it} + \varepsilon_{it}$ and $y_{it}^{\text{gap}} = \tilde{y}_{it}^{\text{gap}} + \varepsilon_{it}$ with $(\tilde{\pi}, \tilde{y})$ solving (2)–(3), so every cross-sectional variance below carries the additive constant $\text{Var}_i(\varepsilon_{it}) = \sigma_\varepsilon^2 \geq 0$, which raises every dispersion floor and weakens none; **(D6)** the inflation target $\pi^* \geq 0$ is a fixed exogenous scalar and the initial cross-sectional mean expectation is non-degenerate, $\bar{\xi}(0) \neq \pi^*$; **(D7)** the cross-sectional distribution of the structural innovations $\theta_i - \bar{\theta}$ (equivalently, of $r_i^* - r^*$) is symmetric around zero. An admissible mechanism $M \in \mathfrak{M}(\mathfrak{D})$ is any measurable expectation-formation protocol in $\mathcal{E}$ subject to: (A) *Capacity Compliance*: Information extraction operates subject to $h_{\text{agg}}(F) \leq h_a$. (B) *Dimensional Restriction*: The monetary instrument is a scalar mapping $R_t : \mathbb{R}^K \to \mathbb{R}$; agents resolve signal extraction via minimum-variance Bayesian optimization against heterogeneous local priors under bounded individual

filtration Ω_{it}. Condition (B) holds for any measurable $R_t : \mathbb{R}^K \to \mathbb{R}$, linear or non-linear; dimensional rank-deficiency is a consequence of the scalar codomain ($N = 1$), not of the functional form. (C) *Structural Consistency*: Every mechanism absorbs the physical transmission primitives $(\kappa_0, \sigma_0, h_a, K, c)$ as invariant environmental parameters, independently of the expectation-formation protocol.

The mathematical domain $\mathfrak{M}(\mathfrak{D})$ bounds the baseline monetary transmission geometry exclusively to operational algorithms projecting scalar outcomes $N = 1$; multi-instrument policy environments parameterized by continuous vector spaces define structurally independent geometric manifolds subject to distinct measurement evaluation.

SEC-1(U) — Universal within $\mathfrak{D}$. The admissible class $\mathfrak{D}$ restricts signal structures to regular exponential families with conjugate priors for which the posterior variance $f(\tau)$ is a strictly decreasing, strictly convex function of total precision τ; this class covers every signal structure in the rational inattention literature (Sims, 2003; Mackowiak and Wiederholt, 2009; Cover and Thomas, 1991). Within this class, $\partial^2 V_{ik}/(\partial h_{ik}\, \partial h_p) \leq 0$ for all i, k; under Gaussian signals, the inequality is strict.

SEC-1(B) — Benchmark threshold. Under Assumption 3.7 and Lemma 3.5, geometric degradation reaches the capacity-constraint release boundary at the aggregate transmission ceiling $F_1 = (1/\gamma_I)\sqrt{D^+V(h_a)h_a/(Kc) - 1}$.

SEC-2 (Credibility Extraction Mapping). The admissibility class $\mathfrak{D}$ restricts the spatial credibility assignment algorithm λ_i strictly to a measurable Bayesian mapping structured iteratively against localized unspanned spatial residuals.

SEC-3 (Nonlinear Welfare Aggregation). The central bank welfare objective preserves the strict concavity primitive $u''(\bar{c}, \bar{n}) < 0$ exclusively through spatial integration across the continuous heterogeneous dispersion field. The term $u''(\bar{c}, \bar{n})$

does not enter the CE-failure argument of Lemma 3.9 (Steps (a)–(b)), which operates exclusively through bounded capability in SEC-1 and SEC-2 alone.

SEC-3 establishes an absolute welfare Jensen deficit: strictly concave regional utilities evaluated across heterogeneous inflation realizations deliver strictly lower aggregate welfare than the utility evaluated at the cross-sectional arithmetic mean, by Jensen's inequality applied to u.

The three bilateral impossibilities — CE–PE (SEC-1, SEC-2), PE–US (SEC-2), CE–US (SEC-3) — operate structurally independently: each is forced by a distinct primary mechanism — instrument dimensional scarcity, topological channel capacity limit, and nonlinear aggregation — whose removal eliminates that bilateral's specific collision while leaving the other two active.

Proposition 2.1 (Structural Necessity under Scalar Architecture)**.** *Within the scalar-instrument domain $\mathfrak{D}$, the foundational primitives (P1)–(P3) and domain constraints (D1)–(D7) enforce exact structural limits across any admissible mechanism in $\mathcal{E}(\mathfrak{D})$: (i) the SEC-1(U) cross-partial evaluates strictly negative uniformly across any non-degenerate regular exponential probability mapping; (ii) finite scalar capacity (P1), structural heterogeneity (P2), rank-deficiency (D3), and spatial sequential symmetries (D5, D7) mathematically bound spatial credibility dispersion strictly positive and guarantee its continuous expansion $d\,\mathrm{Var}_i(\lambda_i|F)/dF > 0$ at every finite t; (iii) strict regional structural concavity (P3) enforces a strictly negative Jensen welfare wedge $W_J < 0$ continuously across any positive spatial expectation variance.*

Proof. (i): Define $V_{ik} \equiv \mathrm{Var}(x|p_t) - \mathrm{Var}(x|p_t, s_{ik})$ as the marginal variance reduction from private signal s_{ik} conditional on public signal p_t. For any regular exponential family with conjugate prior, the posterior variance $f(\tau)$ is strictly decreasing and

strictly convex in total precision $\tau \equiv \tau_0 + h_p + h_{ik}$ (Cover and Thomas, 1991): $f'(\tau) < 0$ and $f''(\tau) > 0$ for all τ in the admissible range. The regular exponential family guarantees exact twice-differentiability (C^2) of the posterior variance function over total precision. Expressing the conditional variance reduction as $V_{ik} = f(\tau_0 + h_p) - f(\tau_0 + h_p + h_{ik})$ executes the strict cross-partial derivative:

$$\frac{\partial^2 V_{ik}}{\partial h_{ik}\, \partial h_p} = -f''(\tau_0 + h_p + h_{ik}) \leq 0,$$

with strict inequality wherever $f'' > 0$, which holds in every non-degenerate member of the admissible class. SEC-1(U) holds across the full admissible class $\mathfrak{D}$ (Sims, 2003; Mackowiak and Wiederholt, 2009). Under Gaussian signals, the posterior variance from the public signal alone is $(h_0 + h_p)^{-1}$; dual-filtration extraction yields $(h_0 + h_{ik} + h_p)^{-1}$. The conditional value is therefore strictly positive: $V_{ik} = (h_0 + h_p)^{-1} - (h_0 + h_{ik} + h_p)^{-1} > 0$, a strictly monotone function of $I(s_{ik}; x|p_t)$.

The derivative $\partial V_{ik}/\partial h_p = -(h_0 + h_p)^{-2} + (h_0 + h_{ik} + h_p)^{-2} < 0$ for all $h_{ik} > 0$. Expanding aggregate precision compresses the marginal contribution of the private signal, as the DPI requires. The cross-partial $\partial^2 V_{ik}/\partial h_{ik} \partial h_p = -2(h_0 + h_{ik} + h_p)^{-3} < 0$ confirms SEC-1(U) strictly over the Gaussian optimizing subclass.

(ii): The performance residual e_{it}, embedded within Ω_{it} as a zero-marginal-cost signal dimension, enters the Bayesian credibility extraction problem without consuming finite bandwidth (P1); bounded capacity restricts reconstruction of the unconfounded spatial state θ_i from the confounded aggregate, sealing credibility dispersion $\mathrm{Var}_i(\lambda_i) > 0$ at every finite t across any admissible mechanism.

Let $q_i \equiv r_i^* - r^*$ and $\hat{q}_i(t)$ its minimum-mean-squared-error recovery from the bank's finite-dimensional aggregate information at date t, and set $\varphi_i(t) \equiv q_i - \hat{q}_i(t)$. Fundamental heterogeneity (P2) establishes $\mathrm{Var}_i(q_i) = F^2 > 0$. Write $q_i = \bar{q} + \zeta_i$

with $\int_0^1 \zeta_i \, di = 0$; by Assumption 2.1, $\{\zeta_i\}$ is drawn independently across regions and independent of the aggregate probability space, so any finite collection of bounded-weight aggregate readings $S_k = \int_0^1 w_k(i) q_i \, di + u_k$, $k = 1, \ldots, n_t < \infty$, satisfies $\int_0^1 w_k(i) \zeta_i \, di = 0$ almost surely by the Exact Law of Large Numbers (Judd, 1985): every such reading is a function of $\bar{q}$ and noise alone, carrying no information about any individual ζ_i, so $\hat{q}_i(t)$ is common across i up to a Lebesgue-null exactly-resolved set, and $\varphi_i(t) = \zeta_i$ almost everywhere. This is what independence across regions buys, and it is why the restriction is load-bearing: a heterogeneity profile generated by a single common factor loading differently across i — which is exactly what Assumption 2.1 excludes, since it is not idiosyncratic — could be read off exactly from one aggregate statistic even at finite capacity, and finite capacity alone would then not force any tracking error at all. Fundamental heterogeneity (P2) and this independence jointly guarantee the strictly positive cross-sectional tracking error $\mathrm{Var}_i(\varphi_i(t)) = \mathrm{Var}_i(\zeta_i) = F^2 > 0$ at every finite t. Substituting this localized differential into the aggregate observable isolates the structurally unspanned residual:

$$\breve{e}_{it} \equiv e_{it} - \psi(\hat{\mu}_i(t) - \bar{\mu}(t)) = \psi \varphi_i(t) - \delta_t(\theta_{CB}) + \varepsilon_{it}.$$

The term ε_{it} is a structural measurement innovation; the variance floor $\mathrm{Var}_i(\varphi_i(t)) > 0$ established in the following paragraph does not invoke any distributional property of ε_{it} and holds across all measurable innovation sequences. This confounds the local unspanned characteristic $\psi \varphi_i(t)$ with the scalar projection $\delta_t(\theta_{CB})$.

Finite capacity $h_a < \infty$ (P1) restricts the local estimator to a finite-rank projection onto a scalar signal permanently confounded by the central-bank policy type θ_{CB}. Extracting the cross-sectional mean space demands geometric inversion of θ_{CB} directly from the scalar residual. Dimensional rank-deficiency (D3) precludes this inversion;

unspanned spatial coordinates persist orthogonal to the aggregate policy vector, consistent with the tracking floor already established.

Let $\lambda_i = \Lambda(\breve{e}_{it})$ denote the Bayesian credibility-assignment operator. Under condition (B), the residual $\breve{e}_{it} = \psi\varphi_i(t) - \delta_t(\theta_{CB}) + \varepsilon_{it}$ satisfies the Monotone Likelihood Ratio Property: by conditions (D5) and (D7), the cross-sectional distribution of $\breve{e}_{it}$ belongs to a location-scale family indexed by the unspanned natural-rate deviation $r_i^* - r^*$, and the likelihood ratio of high versus low central-bank type is strictly monotone in the residual value. Writing the marginal density of $\breve{e}$ as $m(\breve{e}) \equiv \int f(\breve{e} \mid z)\, d\nu(z)$, Bayes' rule gives posterior densities in ratio $d\nu(\cdot \mid \breve{e}_2)/d\nu(\cdot \mid \breve{e}_1)(z) = [f(\breve{e}_2 \mid z)/f(\breve{e}_1 \mid z)] \cdot [m(\breve{e}_1)/m(\breve{e}_2)]$; the second factor does not depend on z, so strict MLRP makes this ratio strictly increasing in z whenever $\breve{e}_2 > \breve{e}_1$, giving strict first-order stochastic dominance of the posterior at the higher signal. Since $\lambda_i = \Lambda(\breve{e}_{it}) = \mathbb{E}[q_i \mid \breve{e}_{it}]$ for a credibility state q_i strictly increasing in z, strict FOSD delivers $\Lambda(\breve{e}_{it})$ strictly increasing in $\breve{e}_{it}$ for any Bayes-optimal Λ preserving this ratio structure (Cover and Thomas, 1991); the result holds for every signal family preserving this ratio ordering, with the Gaussian benchmark as one instance.

The cross-sectional residual $\breve{e}_{it}$ constitutes a location-scale probability family parameterized strictly by spatial divergence F. Expanding F dilates the structural distribution without bound, enforcing strict convex dispersive order across the F-indexed measurement support. Against this geometric dilation, the optimal monotone Bayesian operator Λ migrates output probability mass strictly toward the boundary tails; this continuous structural dispersion forces cross-sectional credibility variance to expand, establishing $d\,\mathrm{Var}_i(\lambda_i|F)/dF > 0$ at every finite t.

(iii): The universal aggregation failure follows from decoupling the spatial policy objective from the informational learning problem. The US criterion requires

$\mathrm{Var}_i(\pi_i - \pi^*) \leq \eta$, imposing a strictly convex quadratic spatial loss.

The Jensen wedge $W_J \equiv \frac{1}{2}u''(\bar{c}, \bar{n})\,\mathrm{Var}_i(\mathbb{E}_i[x])$ bounds the sigma-algebra generated by the agents' learning problem. Regional utility is strictly concave ($u''(\bar{c}, \bar{n}) < 0$, Primitive P3); optimal extraction therefore respects convex information degradation, a consequence of the DPI and capacity limits. Jensen's inequality within the informational filtration forces an extraction deficit proportional to $W_J \neq 0$ whenever $\mathrm{Var}_i(\mathbb{E}_i[x]) > 0$.

Under Assumption 3.2 with $f_i = 0$, Lemma 3.2 gives $\pi_i - \bar{\pi} = a_\pi q_i$ directly, so

$$\mathrm{Var}_i(\pi_i - \pi^*) = a_\pi^2\,\mathrm{Var}_i(r_i^*) = a_\pi^2 F^2 > 0 \tag{6}$$

for every $F > 0$, with no auxiliary comonotonicity argument required: the natural-rate deviation is, under this benchmark, the exact and only source of cross-sectional inflation variation. Two parallel conclusions follow. *(iii-a) Stabilization floor.* The uniform stabilization criterion η inherits the physical floor $a_\pi^2 F^2 > 0$ from instrument scarcity and expectation dispersion, independently of the welfare aggregator. *(iii-b) Welfare deficit.* Jensen's inequality operating over the strictly concave regional utility (Primitive P3) forces aggregate welfare evaluation strictly below the symmetric coordinate center. Evaluating the exact aggregation integral establishes a structural welfare gap $W_J < 0$ scaling strictly proportionate to the unspanned spatial dispersion $a_\pi^2 F^2$, independently of the physical stabilization floor. ■

3. Structural Foundations of the Impossibility

3.1 Information Coordination Failure

Differencing (3) across regions determines the real block directly: $y_i^{\text{gap}} - \bar{y}^{\text{gap}} = \sigma_0(r_i^* - r^*) + \sigma_0 f_i + h_i^+$, the forward inflation forecast entering with a positive sign because a higher expected price level lowers the real financing cost. What the current inflation deviation $\pi_i - \bar{\pi}$ equals depends on how (2) is closed, and MTP-3's own timing forces a discount-rate factor into that closure that a purely contemporaneous reading of the Phillips curve would miss.

Assumption 3.1 (Instrument Common Knowledge — Domain Restriction)**.** Regional date-$(t-1)$ forecasts of the common policy instrument are undispersed: $\mathbb{E}_{i,t-1}[R_t] = \int_0^1 \mathbb{E}_{j,t-1}[R_t]\, dj$. Regions may misforecast the instrument; they may not disagree about it. This is a restriction on information, not a consequence of the scalar instrument.

For $q_i \equiv r_i^* - r^*$, define the three deviation objects that carry the dynamic cross-section:

$$h_i \equiv \mathbb{E}_{i,t-1}[y_{it}^{\text{gap}}] - \int_0^1 \mathbb{E}_{j,t-1}[y_{jt}^{\text{gap}}]\, dj, \quad h_i^+ \equiv \mathbb{E}_{i,t-1}[y_{i,t+1}^{\text{gap}}] - \int_0^1 \mathbb{E}_{j,t-1}[y_{j,t+1}^{\text{gap}}]\, dj,$$

$$f_i \equiv \mathbb{E}_{i,t-1}[\pi_{i,t+1}] - \int_0^1 \mathbb{E}_{j,t-1}[\pi_{j,t+1}]\, dj.$$

Lemma 3.1 (Dynamic Cross-Sectional Inflation)**.** *Under* (2), (3), *MTP-3, Assumption 2.6, and Assumption 3.1, on* $[0, F_1)$*,*

$$(1-\beta)(\pi_{it} - \bar{\pi}_t) = \kappa_0 \sigma_0 q_i + \kappa_0 h_i^+ + \kappa_0 \sigma_0 f_i. \tag{7}$$

No restriction on h_i^+ or f_i is imposed.

Proof. MTP-3 makes π_{it} measurable with respect to $\Omega_{i,t-1}$, so $\mathbb{E}_{i,t-1}[\pi_{it}] = \pi_{it}$ identically — this is idempotence of conditional expectation on a measurable function, not a further assumption — and (2) reads $\pi_{it} = \beta\pi_{it} + \kappa_0 \mathbb{E}_{i,t-1}[y_{it}^{\text{gap}}]$. Collecting terms and differencing across regions,

$$(1-\beta)(\pi_{it} - \bar{\pi}_t) = \kappa_0 h_i. \tag{8}$$

The discount factor enters exactly here: predetermination moves the $\beta\pi_{it}$ term to the left rather than cancelling it, so the Phillips block transmits the forecast gap at rate $\kappa_0/(1-\beta)$, not κ_0. Since $\Omega_{i,t-1} \subset \Omega_{it}$, taking $\mathbb{E}_{i,t-1}[\cdot]$ of (3) and applying the tower property gives

$$\mathbb{E}_{i,t-1}[y_{it}^{\text{gap}}] = \mathbb{E}_{i,t-1}[y_{i,t+1}^{\text{gap}}] - \sigma_0\big(\mathbb{E}_{i,t-1}[R_t] - \mathbb{E}_{i,t-1}[\pi_{i,t+1}] - r_i^*\big).$$

Differencing across regions and applying Assumption 3.1 to delete the instrument-forecast term, $h_i = h_i^+ + \sigma_0 f_i + \sigma_0 q_i$; substituting into (7) delivers (6). ■

Assumption 3.2 (Forward-Gap Benchmark — Domain Restriction)**.** The future output-gap channel carries no cross-sectional dispersion, $h_i^+ = 0$. No restriction is imposed on f_i.

Lemma 3.2 (Forward-Benchmark Inflation Loading)**.** *Under Assumption 3.2,* $(1-\beta)(\pi_{it} - \bar{\pi}_t) = \kappa_0\sigma_0(q_i + f_i)$*. If future inflation forecasts are undispersed,* $f_i = 0$*, then*

$$\pi_{it} - \bar{\pi}_t = a_\pi\, q_i, \qquad a_\pi \equiv \frac{\kappa_0\sigma_0}{1-\beta}, \tag{9}$$

with no further condition on $\kappa_0, \sigma_0, \beta$ *required.*

Proof. Set $h_i^+ = 0$ in (6); if also $f_i = 0$, divide by $1 - \beta > 0$. ■

Lemma 3.3 (Exact Product-Variance Identity)**.** *Let $X_i = \bar{X} + U_i$ and $Y_i = \bar{Y} + V_i$ with $\mathbb{E}_i[U_i] = \mathbb{E}_i[V_i] = 0$ and all fourth moments finite. Then*

$$\begin{aligned}\mathrm{Var}_i(X_iY_i) = \bar{X}^2\,\mathrm{Var}_i(Y_i) + \bar{Y}^2\,\mathrm{Var}_i(X_i) + \mathrm{Var}_i(U_iV_i) \\ + 2\bar{X}\bar{Y}\ \mathrm{Cov}_i(X_i, Y_i) + 2\bar{X}\,\mathbb{E}_i[U_iV_i^2] + 2\bar{Y}\,\mathbb{E}_i[U_i^2V_i].\end{aligned} \tag{10}$$

The identity is exact and carries no sign implication: every sign used downstream must be supplied either by the symmetry route (P-odd with D-i and D-ii, Lemma 3.8) or, for nonsymmetric cross-sections, by the fallback domain of Assumption 3.4.

Proof. Expand $X_iY_i = \bar{X}\bar{Y} + \bar{X}V_i + \bar{Y}U_i + U_iV_i$, subtract the mean $\bar{X}\bar{Y} + \mathbb{E}_i[U_iV_i]$, square, and integrate; the cross terms are exactly the three displayed covariance and third-moment terms. ■

Assumption 3.3 (Full Anchoring — Benchmark Restriction)**.** Both forward channels are cross-sectionally undispersed: $h_i^+ = 0$ and $f_i = 0$, hence $w_i \equiv f_i + h_i^+/\sigma_0 = 0$. Assumption 3.3 strictly strengthens Assumption 3.2: the latter restricts only h_i^+ and carries no implication for f_i, so $f_i = 0$ is never inferred from Assumption 3.2. Every result stated under Assumption 3.3 is benchmark-specific and is never invoked in the proof of Theorem 4.1, of Lemma 3.5, or of Corollary 3.3.

Assumption 3.4 (Residual Covariance Sign Domain — Fallback Restriction)**.** With $X_i = \bar{X} + U_i$ and $Y_i = \bar{Y} + V_i$ as in Lemma 3.3, the third-order residual terms satisfy

$$\mathbb{E}_i[U_iV_i^2] \geq 0, \qquad \mathbb{E}_i[U_i^2V_i] \cdot \bar{Y} \geq 0, \qquad \mathrm{Cov}_i[X_i, Y_i] \cdot \bar{X}\bar{Y} \leq 0.$$

Assumption 3.5 (Stationary-Cross-Section Benchmark — Domain Restriction)**.** For the static benchmark only, $h_i^+ = 0$ and $f_i = \pi_{it} - \bar{\pi}_t$: next-period inflation forecasts are identified with the current cross-sectional inflation deviation. This is a closure condition, not a consequence of MTP-3.

Assumption 3.6 (Gap Predetermination — Domain Restriction)**.** For the realized-gap representation of the stationary benchmark, $y_{it}^{\text{gap}} \in \Omega_{i,t-1}$, so $\mathbb{E}_{i,t-1}[y_{it}^{\text{gap}}] = y_{it}^{\text{gap}}$.

Lemma 3.4 (Static AS-AD Closure)**.** *Under Lemma 3.1, Assumptions 3.5 and 3.6, the stationary-cross-section benchmark satisfies*

$$y_i^{\text{gap}} - \bar{y}^{\text{gap}} = \gamma_c \, q_i, \qquad \gamma_c \equiv \frac{(1-\beta)\sigma_0}{(1-\beta) - \kappa_0\sigma_0}, \qquad \pi_i - \bar{\pi} = \frac{\kappa_0\sigma_0}{(1-\beta) - \kappa_0\sigma_0} q_i, \quad (11)$$

unique and with $\gamma_c > 0$ *if and only if*

$$\textbf{\textit{(SI)}} \qquad \kappa_0\sigma_0 < 1 - \beta. \quad (12)$$

Under (11)*,* $\text{Var}_i(y_i^{\text{gap}}) = \gamma_c^2 F^2 > 0$ *for every* $F > 0$*, and the scalar instrument leaves* $\gamma_c^2 F^2$ *unspanned.*

Proof. Under Assumption 3.5, (6) becomes $(1-\beta)(\pi_i - \bar{\pi}) = \kappa_0\sigma_0 q_i + \kappa_0\sigma_0(\pi_i - \bar{\pi})$; the second term's sign is inherited from the plus sign established in Lemma 3.1's proof, so the static block is self-reinforcing, not self-correcting. Solving, $(\pi_i - \bar{\pi})[(1-\beta) - \kappa_0\sigma_0] = \kappa_0\sigma_0 q_i$, invertible with a positive coefficient exactly under (11). Assumption 3.6 and the date-$(t-1)$ forecast of (3) give $y_i^{\text{gap}} - \bar{y}^{\text{gap}} = \sigma_0 q_i + \sigma_0(\pi_i - \bar{\pi})$; substituting delivers (10). ■

Remark 3.1 (The Static AS-AD Coefficient γ_A). An alternative static block, $\pi_i - \bar{\pi} = \kappa_0(y_i^{\text{gap}} - \bar{y}^{\text{gap}})$, $y_i^{\text{gap}} - \bar{y}^{\text{gap}} = \sigma_0 q_i - \sigma_0(\pi_i - \bar{\pi})$, gives the coefficient $\gamma_A \equiv \sigma_0/(1+\kappa_0\sigma_0)$. That block omits the discount factor and assigns current inflation a negative loading in the real financing-cost relation; it is internally coherent but is not the closure of (2)–(3) under MTP-3, and is not used elsewhere in this paper.

Assumption 3.7 (Extraction Payoff Structure — Benchmark)**.** The information payoff to aggregate attention takes the separable form $\rho(h_{\text{agg}}, F)^2 = V(h_{\text{agg}}) \cdot \rho_0(F)^2$, where $\rho_0(F)^2 = 1/(1+\gamma_I^2 F^2)$ captures the fragmentation-driven decay in signal fidelity and $\gamma_I \equiv \sigma_0/(1+\kappa_0\sigma_0) > 0$ is a benchmark transmission-sensitivity scalar specified by this information-payoff environment, distinct from the dynamic inflation loading a_π of Lemma 3.2. The capacity multiplier $V : [0, h_a] \to [0, 1]$ satisfies: (i) $V(0) = 0$; (ii) V is strictly increasing and strictly concave; (iii) $D^+V(h_a) > 0$; (iv) $\lim_{h\to 0^+} V'(h) = \infty$. The power-law specification $V(h) = h^{\beta_\rho}$ with $\beta_\rho \in (0, 1)$ satisfies all four conditions: $V'(h) = \beta_\rho h^{\beta_\rho - 1} \to \infty$ as $h \to 0^+$. Lemma 3.5 holds for any V satisfying (i)–(iv), with F_1 determined by the boundary condition $D^+V(h_a)(h_a/K)\rho_0(F_1)^2 = c$, where the strictly positive scalar $c > 0$ defines the constant marginal cost of spatial extraction.

Lemma 3.5 (Aggregate Transmission Ceiling)**.** *Under Assumption 3.7 and non-degeneracy $D^+V(h_a)h_a/(Kc) > 1$, a unique positive threshold*

$$F_1 \equiv \frac{1}{\gamma_I}\sqrt{\frac{D^+V(h_a)\, h_a}{Kc} - 1} > 0 \tag{13}$$

exists; above F_1 the capacity constraint releases and $\kappa(F)\sigma(F) < \delta$ for any $\delta > 0$ above a finite $F(\delta)$, with $\lim_{F\to\infty} \kappa(F)\sigma(F) = 0$.

Proof. From Assumption 3.7, the extraction payoff is $\rho(h_{\text{agg}}, F)^2 = V(h_{\text{agg}}) \cdot \rho_0(F)^2$ where $\rho_0(F)^2 = 1/(1 + \gamma_I^2 F^2)$. Strict concavity of V delivers a strictly decreasing

superdifferential; let $D^+V(h)$ denote the upper right Dini superdifferential of V at h. The superdifferential first-order condition for an interior maximum is $D^+V(h_{\text{agg}})(h_a/K)\rho_0(F)^2 = c$. For $F < F_1$, the viability condition $D^+V(h_a)(h_a/K)\rho_0(F)^2 > c$ holds; strict concavity forces the objective to increase in h throughout $[0, h_a)$:

$$D^+\Pi(h) \equiv D^+V(h)(h_a/K)\rho_0(F)^2 - c > 0 \quad \text{for all } h \in [0, h_a).$$

Because V is strictly concave, the Dini superdifferential D^+V is strictly decreasing on $(0, h_a]$. Condition (iv) of Assumption 3.7 forces $D^+V(h) \to \infty$ as $h \to 0^+$; the unique global maximum on $[0, h_a]$ therefore occurs at the boundary corner $h_{\text{agg}}(F) = h_a$. At the defined capacity threshold, $D^+V(h_a)(h_a/K)\rho_0(F_1)^2 = c$ holds with strict equality. For all $F > F_1$, the inequality $D^+V(h_a)(h_a/K)\rho_0(F)^2 < c$ eliminates the boundary corner from optimality. Condition (iv) of Assumption 3.7 supplies $D^+V(0^+) = \infty$, so $h = 0$ is also not optimal. Because D^+V is strictly decreasing on $(0, h_a]$ and $\rho_0(F)^2$ is strictly decreasing in F, the first-order condition $D^+V(h_{\text{agg}}(F))(h_a/K)\rho_0(F)^2 = c$ implicitly defines $h_{\text{agg}}(F)$ as a strictly decreasing function of F. Hence $h_{\text{agg}}(F) \in (0, h_a)$ for all $F > F_1$, with $h_{\text{agg}}(F) \to 0^+$ as $F \to \infty$. ■

Corollary 3.1 (Policy Transmission Collapse)**.** *For $F > F_1$, the capacity constraint releases and $h_{\text{agg}}(F)$ is strictly interior to $(0, h_a)$, decaying continuously to zero as $F \to \infty$. Under Assumption 2.6, $\kappa(F)\sigma(F) > 0$ for every finite $F > F_1$, but for any transmission tolerance $\delta > 0$ there exists a finite threshold $F(\delta)$ such that $\kappa(F)\sigma(F) < \delta$ for all $F > F(\delta)$.*

Lemma 3.6 (Maximal Capacity Saturation)**.** *For $F < F_1$, the extraction condition dictates maximal capacity utilization $h_{\text{agg}}(F) = h_a$. Under Assumption 2.6, this stabilizes the global baseline: $\kappa(F)\sigma(F) = \kappa_0\sigma_0$ on $[0, F_1)$.*

Proof. Lemma 3.5 establishes $h_{\text{agg}}(F) = h_a$ on $[0, F_1)$ under the viability condition. Assumption 2.6 then delivers $\kappa(F)\sigma(F) = \kappa_0\sigma_0$ on $[0, F_1)$ and $\kappa(F)\sigma(F) > 0$ for every finite $F > F_1$. ■

Corollary 3.2 (Transmission Class)**.** *The PE impossibility boundary holds globally for any continuous mapping $\kappa(F)$ satisfying $\kappa(F)\sigma(F) > 0$ for all finite F and $\lim_{F\to\infty} \kappa(F)\sigma(F) = 0$; Assumption 2.6 is one member of this generic class and no architecture inside $\mathfrak{M}(\mathfrak{D})$ exits it.*

3.2 Credibility Heterogeneity

Channel (F), universal across $\mathfrak{M}(\mathfrak{D})$: Bounded extraction capacity $h_a < \infty$ (P1) and dimensional rank-deficiency (D1–D3) force strictly positive cross-sectional tracking dispersion $\text{Var}_i(\varphi_i(t)) > 0$ at every finite t, generating $\text{Var}_i(\lambda_i) > 0$ independently of any drift assumption. Channel (S1), requiring Assumption 2.3: Under A-RW, constant-gain extraction dictates the optimal filtration; geometric non-concentration bounds sustain $\text{Var}_i(\lambda_i) > 0$ in the limit distribution, securing $C_\lambda^{\text{stat}} > 0$. Channel (S2), requiring $\sigma_\varepsilon^2 > 0$: Irreducible dimensional measurement drift limits expectation convergence, establishing $\text{Var}_i(\xi_i) > 0$ globally asymptotically.

Lemma 3.7 (Credibility Divergence)**.** *For any $F > 0$ and any mechanism in $\mathfrak{M}(\mathfrak{D})$, $\text{Var}_i(\lambda_i) = g(F, t) > 0$ at every finite horizon t. The dispersion mapping $g(F, t)$ is Lipschitz continuous on $[0, F_1)$ with finite constant $L_g < \infty$.*

Proof. At the homogeneous limit $F = 0$, the structural boundary condition $\text{Var}_i(r_i^*) = 0$ dictates $r_i^* = r^*$ across the entire spatial measure. This total collapse of structural variation to the scalar mean nullifies the persistent capacity-constrained tracking bias ($\varphi_i(t) = 0$). The Bayesian updating operator maps strictly relative to the persistent structural differential; integrating out the spatially orthogonal, transient idiosyncratic

noise (ε_{it}) isolates the structurally measurable residual to the common aggregate term $-\delta_t(\theta_{CB})$. The consequent credibility assignment λ_i projects uniformly against this invariant aggregate scalar across space, algebraically generating the absolute origin boundary $g(0,t) \equiv \mathrm{Var}_i(\lambda_i) = 0$.

For any $F > 0$, non-null dimensional binding-frequency deviations manifest cross-sectionally. Under the Bayesian extraction matrix of condition (B), $\breve{e}_{it}$ is non-degenerate across i and Λ is strictly increasing on its support (Proposition 2.1(ii)), hence injective there, so $\lambda_i = \Lambda(\breve{e}_{it})$ takes at least two values on sets of positive measure and $g(F,t) > 0$. This part uses only injectivity and holds for any monotone Bayes-optimal Λ.

Strict monotonicity $\partial g/\partial F > 0$ requires more: boundedness $\lambda_i \in [0,1]$ forces Λ to saturate, and a saturating Λ can *compress* dispersion as F rises if the cross-section is not symmetric — raising F can push every region into the same saturated branch, lowering rather than raising $\mathrm{Var}_i(\lambda_i)$ — so the claim is not free of restriction on the credibility rule. Under Assumption 3.8, writing $\Lambda = \bar{\lambda}(F) + A$ with A the odd part, symmetry of the cross-section makes $g(F,t) = \mathbb{E}_i[A(\breve{e}_{it})^2]$ exactly (the even part contributing only to the mean), and A strictly increasing and odd makes $A(\breve{e}_{it})^2$ strictly increasing in $|\breve{e}_{it}|$; since raising F strictly increases the scale of $\breve{e}_{it}$ pointwise wherever it is non-zero, $g(F,t)$ is strictly increasing. This is the case used throughout the paper; the qualifier is carried explicitly rather than claimed to be free.

$\lambda_i \in [0,1]$ is absolutely bounded, delivering $\|\lambda_i - \bar{\lambda}\|_{L^2} \leq 1/2$. The mapping $F \mapsto g(F,t) = \mathrm{Var}_i(\lambda_i)$ is absolutely continuous on $[0, F_1)$: the dominated convergence theorem applied to the bounded integrand $(\lambda_i - \bar{\lambda})^2$ transfers continuity from the scale-parameterized distribution of $\breve{e}_{it}$ directly to $g(F,t)$. The primitives $(\psi, \|\bar{\theta}\|, \sigma_\varepsilon^2)$ are bounded by domain construction; the spatial first difference quotient of $g(F,t)$ is bounded by the product of these finite scalars and the Lipschitz constant of Λ

on $[0,1]$, delivering $L_g < \infty$ on $[0, F_1)$ independently of pointwise differentiability of Λ. ■

Remark 3.2 (Three-Horizon Scope). Three temporal channels operate independently. Channel (F): $A(F,t) > 0$ binds at every finite t unconditionally across $\mathfrak{M}(\mathfrak{D})$, without A-RW. Channel (S1): A-RW sustains the unique stationary credibility measure with $C_\lambda^{\text{stat}} > 0$ via Krylov-Bogolyubov; global irreducibility upgrades existence to uniqueness. Channel (S2): Permanent noise $\sigma_\varepsilon^2 > 0$ sustains $\text{Var}_i(\xi_i) > 0$ globally; F_3^* requires the dual limit distributions from Channels (S1) and (S2).

3.3 Expectation Feedback Amplification

Assumption 3.8 (Symmetric Credibility Response — Domain Restriction)**.** The $\Omega_{i,t-1}$-measurable component of credibility carries a conditional mean antisymmetric about the homogeneous point: for Lebesgue-almost every θ in the support of $r_i^* - r^*$,

$$\mathbb{E}[\lambda_i \mid r_i^* - r^* = \theta] + \mathbb{E}[\lambda_i \mid r_i^* - r^* = -\theta] = 2\bar{\lambda}(F).$$

A region's credibility loss from running θ above the mean natural rate equals its credibility gain from running θ below it; the restriction is satisfied by any credibility rule that is itself antisymmetric in the signed deviation, including the affine benchmark of Proposition 5.1.

Lemma 3.8 (Bias Cancellation Under Symmetric Credibility)**.** *Let $U_i = P_i + \varepsilon_i$ with $P_i \equiv \mathbb{E}[U_i \mid \Omega_{i,t-1}]$ and $\mathbb{E}[\varepsilon_i \mid \Omega_{i,t-1}] = 0$. Suppose there is a cross-sectional index $z_i \equiv r_i^* - r^*$, symmetric about zero (D7), with $P_i = p(z_i)$ and $V_i = v(z_i)$ both odd (Lemma 3.4 for V_i; Assumption 3.8 for P_i), V_i measurable with respect to $\Omega_{i,t-1}$, and $\mathbb{E}[\varepsilon_i^2 \mid z_i = z]$ even in z. Then $\mathbb{E}_i[U_i V_i^2] = \mathbb{E}_i[U_i^2 V_i] = 0$.*

Proof. $\mathbb{E}_i[U_i V_i^2] = \mathbb{E}_i[P_i V_i^2] + \mathbb{E}_i[\varepsilon_i V_i^2]$. $P_i V_i^2 = p(z_i)v(z_i)^2$ is odd (odd times even), so its expectation vanishes under the symmetric cross-section; $\mathbb{E}_i[\varepsilon_i V_i^2] = \mathbb{E}_i[V_i^2\,\mathbb{E}[\varepsilon_i \mid \Omega_{i,t-1}]] = 0$ since V_i is $\Omega_{i,t-1}$-measurable. For the second moment, $\mathbb{E}_i[U_i^2 V_i] = \mathbb{E}_i[P_i^2 V_i] + 2\mathbb{E}_i[P_i \varepsilon_i V_i] + \mathbb{E}_i[\varepsilon_i^2 V_i]$. $P_i^2 V_i$ is odd (even times odd), so it vanishes; $P_i V_i$ is $\Omega_{i,t-1}$-measurable, so $\mathbb{E}_i[P_i \varepsilon_i V_i] = 0$ by the same tower argument; writing $\sigma_\varepsilon^2(z) \equiv \mathbb{E}[\varepsilon_i^2 \mid z_i = z]$, the last term is $\mathbb{E}_i[\sigma_\varepsilon^2(z_i)v(z_i)]$, even times odd, hence zero. ■

Lemma 3.9 (Expectation Feedback: Universal Floor and Eigenvalue Benchmark)**.** *The strict finite-horizon dispersion floor* $A(F,t) = g(F,t)(\pi^* - \bar{\xi}(t))^2 > 0$ *(where* $g(F,t) = \mathrm{Var}_i(\lambda_i) > 0$ *for all* $F > 0$ *and all finite* t *via Lemma 3.7) binds across every admissible mechanism in* $\mathfrak{M}(\mathfrak{D})$ *through Universal Channel (F). An initial non-degenerate prior (*$\bar{\xi}(0) \neq \pi^*$*) subjects the aggregate learning filtration to the finite Shannon capacity constraint (P1). The continuous capacity bound restricts the posterior expectation step-size, structurally guaranteeing the aggregate target gap satisfies* $(\pi^* - \bar{\xi}(t))^2 > 0$ *uniformly across every finite horizon* $t < \infty$*. Operating orthogonally across the spatial dimension, scalar rank-nullity (D1–D3) forces credibility dispersion* $g(F,t) > 0$ *(Lemma 3.7). Multiplying these distinct limits algebraically executes the strictly positive finite-horizon variance floor unconditionally across the admissible class:* $A(F,t) = g(F,t)(\pi^* - \bar{\xi}(t))^2 > 0$*. Bounding the permanent asymptote* $\liminf_{t\to\infty} A(F,t) \geq \underline{A}(F) > 0$ *requires the persistent structural drift of A-RW: mapping optimal filtration against continuous state variance prevents aggregate absorption* $\bar{\xi}(t) \not\to \pi^*$*, permanently securing the infinite-horizon dimensional tracking floor.*

The floor $A(F,t) > 0$ *holds universally across every mechanism in* $\mathfrak{M}(\mathfrak{D})$*; the expectation-divergence coefficient below is benchmark-specific, derived under the*

recursive-anchoring mechanism of Assumption 3.10 and Assumption 3.2 with $f_i = 0$, and under the named condition (D-i)*: $\bar{\pi}(t-1) \leq \pi^*$, which Lemma 3.10 establishes exactly (as $\bar{\pi}(t-1) = 0$) under Assumption 3.9. Under Assumption 3.10 and (D-i), forecast dispersion satisfies $\sigma_F^2(t) \geq \mu_c(F)\sigma_F^2(t-1) + \underline{I}(F)$, where $\sigma_F^2(t) \equiv \mathrm{Var}_i(\mathbb{E}_{i,t-1}[\pi_{i,t+1}])$ is one-step-ahead forecast dispersion measured at $t-1$ — not realized inflation dispersion $V_\pi(t) \equiv \mathrm{Var}_i(\pi_{it})$, which MTP-3 makes identical to $\mathrm{Var}_i(\mathbb{E}_{i,t-1}[\pi_{it}])$ and therefore unusable as an independent propagating state. The expectation-divergence coefficient is:*

$$\mu_c(F) = \left[\frac{\kappa_0 \sigma_0 \, \bar{X}(F)}{1-\beta}\right]^2, \qquad \bar{X}(F) \equiv 1 - \bar{\lambda}^-(F), \tag{14}$$

with $\bar{\lambda}^-(F) \equiv \int_0^1 \mathbb{E}_{i,t-1}[\lambda_i]\, di$ the projected mean credibility. The injection floor satisfies $\underline{I}(F) \geq \mu_c(F)F^2 > 0$ for every $F > 0$: expectation dispersion never vanishes, though divergence requires the separate condition $\mu_c(F) > 1$.

Proof. The scalar projection $\delta_t : \mathbb{R}^K \to \mathbb{R}$ admits no right inverse (D3); no estimator in $\mathfrak{M}(\mathfrak{D})$ recovers the component of θ_{CB} lying in $\ker(\delta_t)$. Every feasible filtration satisfying condition (B) therefore carries strictly positive unspanned residual variance $\mathrm{Var}(\tilde{\theta}_i|\delta_t) \geq \mathrm{Var}(\tilde{\theta}_i)\cdot\dim(\ker(\delta_t))/K > 0$ (D1–D3), independently of the distributional family governing the state; any credibility-consistent weight assignment therefore satisfies $1 - \lambda_i > 0$ for all $i \in [0,1]$ and every finite t.

Extracting a scalar projection from a multidimensional domain enforces orthogonal residual variance across the spatial measure independently of the mechanism. Bounded capacity restricts extraction strictly, securing the localized credibility dispersion floor $g(F,t) > 0$ (Lemma 3.7); finite Shannon capacity (P1) bounds the aggregate posterior updating rate; given the non-degenerate initial prior (D6), this restriction prevents $\bar{\xi}(t)$ from reaching π^* at any finite t, delivering

$(\pi^* - \bar{\xi}(t))^2 > 0$. The infimum satisfies:

$$A(F,t) \equiv \inf_{M \in \mathfrak{M}(\mathfrak{D})} \mathrm{Var}\big(\mathbb{E}_i[\pi_{t+1}] - \mathbb{E}[\pi_{t+1}]\big) > 0, \tag{15}$$

equal to $g(F,t)(\pi^* - \bar{\xi})^2 > 0$, locking the expectation floor independently of functional covariance forms.

Assumption 3.10 states the benchmark mechanism at the filtration t, $\mathbb{E}_{i,t}[\pi_{i,t+1}] = \lambda_i \pi^* + (1-\lambda_i)\pi_i(t-1)$; identity (6) requires it at filtration $t-1$, one step earlier, since $\pi_{i,t+1}$ is $\Omega_{i,t}$-measurable and $\Omega_{i,t-1} \subset \Omega_{i,t}$ is a strict subset. Applying $\mathbb{E}_{i,t-1}[\cdot]$ to the mechanism and using the tower property, $\pi_i(t-1)$ and π^* pass through unchanged (both are $\Omega_{i,t-1}$-measurable) while λ_i is replaced by its $t-1$ projection $\lambda_i^- \equiv \mathbb{E}_{i,t-1}[\lambda_i]$, giving the re-timed relation

$$\mathbb{E}_{i,t-1}[\pi_{i,t+1}] - \pi^* = X_i Y_i, \qquad X_i \equiv 1 - \lambda_i^-, \qquad Y_i \equiv \pi_i(t-1) - \pi^*,$$

hence $\sigma_F^2(t) = \mathrm{Var}_i[X_i Y_i]$, with $\sigma_F^2(t)$, not the t-conditioned object of the original mechanism, the correct propagating state: $\mathbb{E}_{i,t}[\pi_{i,t+1}]$ is itself $\Omega_{i,t}$-measurable by MTP-3 applied one period ahead, so $\mathrm{Var}_i(\mathbb{E}_{i,t}[\pi_{i,t+1}]) = V_\pi(t+1)$ identically — the same self-reference the master identity's retirement of $\sigma_F^2(t)$ already forced, recurring one index later unless the mechanism itself is re-timed. Decompose $X_i = \bar{X} + U_i$, $Y_i = \bar{Y} + V_i$ with $\bar{X} = 1 - \bar{\lambda} > 0$ and $\bar{Y} = \bar{\pi}(t-1) - \pi^*$ (sign unrestricted here — the bound below holds regardless; Lemma 3.10 establishes $\bar{Y} \le 0$ under stated initial conditions, used elsewhere). The exact product-variance identity of Lemma 3.3 is

$$\begin{aligned}\mathrm{Var}[X_i Y_i] = {} & (\bar{X})^2 \,\mathrm{Var}[Y_i] + (\bar{Y})^2 \,\mathrm{Var}[X_i] + \mathrm{Var}[U_i V_i] \\ & + 2\bar{X}\bar{Y}\ \mathrm{Cov}_i[X_i, Y_i] + 2\bar{X}\,\mathbb{E}[U_i V_i^2] + 2\bar{Y}\,\mathbb{E}[U_i^2 V_i].\end{aligned} \tag{16}$$

Lemma 3.4 makes V_i exactly linear, hence odd, in the natural-rate deviation; Assumption 3.8 makes the $\Omega_{i,t-1}$-measurable component of U_i odd in the same variable; Lemma 3.8 then delivers $\mathbb{E}[U_i V_i^2] = \mathbb{E}[U_i^2 V_i] = 0$.

The credibility weight $X_i = 1-\Lambda(\breve{e}_{it})$ depends on $\breve{e}_{it} = \eta_i+\varepsilon_{it}$, where $\eta_i \equiv \psi\varphi_i(t)-\delta_t(\theta_{CB})$ is $\Omega_{i,t-1}$-measurable and ε_{it} is serially independent of $\Omega_{i,t-1}$ with $\mathrm{Var}(\varepsilon_{it}) = \sigma_\varepsilon^2 > 0$ (D5). $Y_i = \pi_i(t-1) - \pi^*$ is itself $\Omega_{i,t-1}$-measurable, hence independent of ε_{it}. By Proposition 2.1(ii), Λ is strictly increasing in $\breve{e}_{it}$, and η_i is non-decreasing in $r_i^* - r^*$, so $X_i = 1-\lambda_i$ is non-increasing in $r_i^* - r^*$ through the η_i channel. Lemma 3.4 gives Y_i strictly increasing in the same variable, so $\mathrm{Cov}_i[X_i, Y_i] \le 0$.

Equality in Cauchy–Schwarz, $\mathrm{Cov}_i[X_i, Y_i] = -\sqrt{\mathrm{Var}[X_i]\,\mathrm{Var}[Y_i]}$, requires $X_i - \bar{X}$ and $Y_i - \bar{Y}$ affinely dependent almost surely. Λ strictly increasing means X_i responds to ε_{it} at every fixed value of η_i; since ε_{it} carries strictly positive variance independent of $\Omega_{i,t-1} \ni Y_i$, X_i carries variation Y_i cannot explain, ruling out exact affine dependence. Cauchy–Schwarz is therefore strict, and substituting it into identity (15) while discarding $\mathrm{Var}[U_i V_i] \ge 0$ delivers

$$\sigma_F^2(t+1) = \mathrm{Var}[X_i Y_i] \ge \left(\bar{X}\sqrt{\mathrm{Var}[Y_i]} - \bar{Y}\sqrt{\mathrm{Var}[X_i]}\right)^2 > 0$$

at every finite t and every $F > 0$, on primitives (D5), MLRP, and $\sigma_\varepsilon^2 > 0$ alone — no parameter configuration excluded.

$\sigma_F^2(t) > 0$ therefore holds at every finite t and every $F > 0$ on primitives (D5), MLRP, and $\sigma_\varepsilon^2 > 0$ alone, independently of (D-i): $\mathrm{Var}_i[U_i V_i] \ge 0$ and $\bar{Y}^2\,\mathrm{Var}_i[X_i] \ge 0$ discard cleanly regardless of sign, and the strict Cauchy–Schwarz argument used above for the universal floor did not invoke $\bar{Y}$'s sign. For the full recursion, the cross term $2\bar{X}\bar{Y}\,\mathrm{Cov}_i[X_i, Y_i]$ is not sign-free: $\bar{X} > 0$ and $\mathrm{Cov}_i[X_i, Y_i] \le 0$ hold unconditionally, so the cross term is non-negative exactly when $\bar{Y} \le 0$, which is (D-i). Under (D-i),

discarding all three remaining non-negative terms and applying Lemma 3.2 under Assumption 3.2 with $f_i = 0$, which gives $\mathrm{Var}_i(\pi_i(t-1)) = a_\pi^2 F^2$ exactly, delivers

$$\sigma_F^2(t) \geq \bar{X}(F)^2 \,\mathrm{Var}_i(Y_i) = \bar{X}(F)^2 \,\mathrm{Var}_i(\pi_i(t-1) - \pi^*) = \bar{X}(F)^2 a_\pi^2 F^2 = \mu_c(F) F^2 > 0,$$

the injection floor of (13). Substituting Proposition 4.1's benchmark identity for $\mathrm{Var}_i(\pi_i(t-1))$ in terms of $\sigma_F^2(t-1)$ instead of the constant $a_\pi^2 F^2$ delivers the full recursion $\sigma_F^2(t) \geq \mu_c(F)\sigma_F^2(t-1) + \mu_c(F)F^2$ with $\mu_c(F)$ exactly as in (13), using (D-i) as established by Lemma 3.10. ■

Assumption 3.9 (Aggregate Forecast Neutrality — Benchmark Restriction)**.** For every date t,

$$\int_0^1 \mathbb{E}_{i,t-1}[y_{it}^{\mathrm{gap}}]\, di = 0. \tag{17}$$

This is a restriction on the information structure, not a consequence of Lemma 2.2: the latter establishes the realized identity $\int_0^1 y_{it}^{\mathrm{gap}}\, di = 0$, an integral of dates-$t$ outcomes, while (16) concerns the integral of region-specific date-$(t-1)$ conditional forecasts. With region-specific filtrations $\Omega_{i,t-1}$, the two objects need not coincide: nothing in the model permits exchanging the region-by-region conditional expectation with the cross-sectional integral, since R_t itself is chosen using date-t information unavailable to any region at $t-1$.

Lemma 3.10 (Below-Target Propagation)**.** *Suppose $\pi^* \geq 0$. Under Assumption 3.9 and MTP-3,*

$$\bar{\pi}(t) = 0 \leq \pi^* \qquad \textit{for every } t. \tag{18}$$

Proof. MTP-3 gives $\mathbb{E}_{i,t-1}[\pi_{it}] = \pi_{it}$, so averaging eq. (2) across regions gives $(1-\beta)\bar{\pi}(t) = \kappa_0 \int_0^1 \mathbb{E}_{i,t-1}[y_{it}^{\mathrm{gap}}]\, di$. Assumption 3.9 sets the right-hand side to zero; since $1-\beta > 0$, $\bar{\pi}(t) = 0$, and $\pi^* \geq 0$ gives the stated inequality. ■

Proposition 3.1 (Finite-Horizon Target Non-Attainment)**.** *Suppose $\pi^* > 0$. Under Assumption 3.9, the benchmark mechanism of Lemma 3.9, and the covariance sign condition of Lemma 4.2,*

$$G(t) > 0 \qquad \text{for every finite } t \geq 0,$$

so $\pi^* \neq \bar{\xi}(t)$ strictly at every finite date — not merely $\bar{\pi}(t) \leq \pi^*$, as in Lemma 3.10, but the aggregate expectation gap itself never closes.

Proof. Lemma 3.10 gives $\bar{\pi}(t-2) = 0$. Averaging the benchmark mechanism, $\bar{\xi}(t) = \bar{\lambda}\pi^* + (1-\bar{\lambda})\cdot 0 + \mathrm{Cov}_i(1-\lambda_i, \pi_i(t-2)) = \bar{\lambda}\pi^* + \mathrm{Cov}_i(1-\lambda_i, \pi_i(t-2))$. By Lemma 4.2's covariance-sign argument, $1-\lambda_i$ is non-increasing in $r_i^* - r^*$ while $\pi_i(t-2)$ is increasing in the same variable, so $\mathrm{Cov}_i(1-\lambda_i, \pi_i(t-2)) \leq 0$ and $\bar{\xi}(t) \leq \bar{\lambda}\pi^*$. Hence $G(t) = \pi^* - \bar{\xi}(t) \geq (1-\bar{\lambda})\pi^* > 0$, since $\pi^* > 0$ and $1-\lambda_i > 0$ for every i (Lemma 3.9's proof) makes $\bar{\lambda} < 1$ strictly. The bound holds directly at every finite t, with no induction and no initial-gap hypothesis required. ■

Lemma 3.11 (A Filtering Benchmark)**.** *Let region i observe $s_{i,t} = r_i^* + \varepsilon_{it}$ with ε_{it} i.i.d. across (i,t), $\mathrm{Var}(\varepsilon_{it}) = \sigma_\varepsilon^2 \in (0,\infty)$, and let $\mathbb{E}_{i,t}[\pi_{i,t+1}] = (1-\omega(F))m(R_t) + \omega(F)s_{it}$, $\omega(F) \equiv F^2/(F^2+\sigma_\varepsilon^2)$, with $m(R_t)$ common across i. This mechanism satisfies (P1)–(P3), (D1)–(D3), (D5), (D7), and SEC-1 through SEC-3: it is admissible. For every finite $F > 0$ and every t,*

$$\sigma_F^2(t) = \frac{F^4}{F^2+\sigma_\varepsilon^2},$$

independent of t and of $\sigma_F^2(t-1)$.

Proof. Admissibility: the mutual information between the forecast and r_i^* is $\frac{1}{2}\log(1+F^2/\sigma_\varepsilon^2) < \infty$, satisfying (P1); (P2) is imposed directly; (P3) and SEC-3

restrict the welfare objective, not the forecast map; $m(R_t)$ common across i leaves (D1)–(D3)'s rank-deficiency untouched; (D5), (D7) hold by the stated i.i.d. symmetric construction; the public/private weights sum to one, satisfying SEC-1; setting $\lambda_i \equiv 1 - \omega_i$ with heterogeneous private precision ordered by the observed residual gives $\Lambda' \geq 0$, satisfying SEC-2.

Since only s_{it} varies with i in the forecast, $\sigma_F^2(t) = \omega(F)^2 \operatorname{Var}_i(s_{it}) = \omega(F)^2(F^2 + \sigma_\varepsilon^2) = F^4/(F^2 + \sigma_\varepsilon^2)$, a function of F and σ_ε^2 alone. ■

Proposition 3.2 (The Anchoring Requirement)**.** *There exist no finite F_3^* and no functions $\mu_c(\cdot) \geq 0$, $I(\cdot) \geq 0$ with $\mu_c(F) > 1$ for $F > F_3^*$ such that $\sigma_F^2(t+1) \geq \mu_c(F)\sigma_F^2(t) + I(F)$ holds for every admissible mechanism.*

Proof. Suppose it held for the mechanism of Lemma 3.11. Fix $F > F_3^*$, write $\mu \equiv \mu_c(F) > 1$. Since $\sigma_F^2(t)$ equals the same constant at every date, the recursion gives $\sigma_F^2(t) \geq \mu\, \sigma_F^2(t) + I(F) > \sigma_F^2(t)$, a contradiction. ■

Proposition 3.2 shows that no universal proof of an explosive coordination threshold can hold across every admissible mechanism: the mechanism of Lemma 3.11 is Bayes-optimal and MLRP-consistent (SEC-2), yet its credibility weight governs the split between the public instrument and the current private signal, not between the target and the region's own lagged realized inflation, and its dispersion never grows. SEC-2 restricts a weight; it does not by itself require the forecast to reference $\pi_i(t-1)$ at all.

Assumption 3.10 (Recursive Anchoring — Domain Restriction)**.** The forecast loads on the region's own lagged realized inflation:

$$\mathbb{E}_{i,t}[\pi_{i,t+1}] = \lambda_i \pi^* + (1 - \lambda_i)\pi_i(t-1), \qquad \lambda_i = \Lambda(\theta_i) \in [0,1], \quad \Lambda' \geq 0,$$

with $\bar{\lambda}(F) \equiv \mathbb{E}_i[\lambda_i] < 1$. Mechanisms satisfying this constitute $\mathfrak{M}_{\mathrm{RF}}(\mathfrak{D}) \subsetneq \mathfrak{M}(\mathfrak{D})$.

Remark 3.3 (This Restriction is Necessary). By Proposition 3.2, no coordination threshold holds on $\mathfrak{M}(\mathfrak{D})$ without some restriction excluding the mechanism of Lemma 3.11. Assumption 3.10 is exactly the property that mechanism lacks: it is the weakest addition to the stated primitives under which the mechanism of Lemma 3.11 is inadmissible, since that mechanism's loading on $\pi_i(t-1)$ is identically zero.

Remark 3.4 (The Two Transmission Coefficients, Resolved). Lemma 3.4 and identity (27) were not computing the same object under different bookkeeping; MTP-3's predetermination of π_{it} forces a $(1-\beta)^{-1}$ loading into any correct closure of (2), which the original derivation of $\gamma_A = \sigma_0/(1+\kappa_0\sigma_0)$ omitted. The corrected dynamic relation, Lemma 3.1, is the source identity; every closed form used elsewhere is one of its closures. Three coefficients now carry distinct, non-competing roles: γ_I (Assumption 3.7) is a benchmark scalar of the information-payoff specification, independent of the dynamic Phillips-curve closure; γ_c (Lemma 3.4) is the stationary-cross-section closure's coefficient, defined only under (11); and $a_\pi = \kappa_0\sigma_0/(1-\beta)$ (Lemma 3.2) is the forward-gap closure's coefficient, defined unconditionally. Theorem 4.1's headline, F_1, and Corollary 4.2 use none of the three. The Coordination Threshold, $\mu_c(F)$, and every quantity built from them use a_π throughout, and are unconditional in $(\kappa_0, \sigma_0, \beta)$ except through the credibility-attenuation condition of Assumption 3.11 and, where the recursion's full amplification rather than only its injection floor is invoked, the disinflationary condition (D-i) of Lemma 4.2 and Lemma 3.9. (D-i) is established exactly, as $\bar{\pi}(t-1) = 0$, by Lemma 3.10 under the explicit benchmark Assumption 3.9; results depending on (D-i) inherit that one clearly labeled restriction and no other.

Theorem 3.1 (Irreducibility of the Cross-Sectional State)**.** *There is no map G with*

$\big(\mathrm{Var}_i(\pi_{it}), \mathrm{Var}_i(y_{it}^{\mathrm{gap}}), \mathrm{Cov}_i(\pi_{it}, y_{it}^{\mathrm{gap}})\big) = G\big($*the same three moments at* $t-1, t-2; F\big)$ *consistent with the exact laws of motion implied by* (2)–(3) *under Assumption 3.10. The cross-sectional distribution of inflation and the output gap carries information no finite set of moments captures: two cross-sections with identical values of* $\mathrm{Var}_i(\pi_{i,t-1})$ *can produce different values of* $\mathrm{Var}_i(y_{it}^{\mathrm{gap}})$*, because the latter's recursion depends on* $\mathrm{Cov}_i(1-\lambda_i, \pi_{i,t-1})$*, which is not a function of* $\mathrm{Var}_i(\pi_{i,t-1})$ *alone.*

Proof. The construction is exact in rational arithmetic; no step is numerical and no quantity is rounded. Fix $\beta = \frac{1}{2}$, $\kappa_0 = 1$, $\sigma_0 = \frac{2}{5}$, so that $a_\pi = \kappa_0\sigma_0/(1-\beta) = \frac{4}{5}$. Write $X_i \equiv 1 - \lambda_i^-$, $Y_i \equiv \pi_i(t-1) - \pi^*$, and $\theta_i \equiv r_i^* - r^*$. The stationary closure of Lemma 3.4 applied region by region gives

$$Y_i = \frac{\sigma_0\,\theta_i}{(1-\beta) - \sigma_0 X_i} = \frac{(2/5)\,\theta_i}{\frac{1}{2} - \frac{2}{5}X_i}, \tag{19}$$

so a region's own credibility weight enters its inflation *level*, not merely its forecast: this is the exact channel through which the joint law of (λ_i^-, θ_i) becomes payoff-relevant. Take $\theta_i \in \{-3, -1, 1, 3\}$ with probability $\frac{1}{4}$ each. The cross-section is symmetric about zero (D7) with

$$m_2 \equiv \mathbb{E}_i[\theta_i^2] = \tfrac{1}{4}(9+1+1+9) = 5, \qquad m_4 \equiv \mathbb{E}_i[\theta_i^4] = \tfrac{1}{4}(81+1+1+81) = 41.$$

Cross-section A sets $\lambda_i^- = 0$ on $|\theta_i| = 1$ and $\lambda_i^- = 1$ on $|\theta_i| = 3$. Cross-section B sets $\lambda_i^- = \frac{1}{2}$ on $|\theta_i| = 3$ and, on $|\theta_i| = 1$, draws $\lambda_i^- \in \{1, \frac{1}{2}, 0\}$ with conditional probabilities $\{\frac{7}{23}, \frac{9}{23}, \frac{7}{23}\}$, which sum to one. Both cross-sections therefore share the entire law of θ_i, hence m_2 and m_4, and both record the same mean credibility and the same realized

inflation dispersion, exactly:

$$\bar{\lambda}^{-}(A) = \bar{\lambda}^{-}(B) = \tfrac{1}{2}, \qquad V_{\pi}(A) = V_{\pi}(B) = \tfrac{272}{25}.$$

They record different one-step-ahead forecast dispersion, exactly:

$$\sigma_F^2(A) = \mathrm{Var}_i(X_iY_i) = 8, \qquad \sigma_F^2(B) = \mathrm{Var}_i(X_iY_i) = \tfrac{104}{23}, \qquad \sigma_F^2(A) - \sigma_F^2(B) = \tfrac{80}{23} > 0.$$

The recorded state $(m_2, m_4, \bar{\lambda}^{-}, V_{\pi})$ is thus common to two admissible cross-sections whose propagating state differs by a strictly positive rational amount. No map G from recorded moments to σ_F^2 exists, and no finite moment list closes the recursion: the obstruction is the joint law of $(\lambda_i^{-}, \theta_i)$, which enters $\mathrm{Var}_i(X_iY_i)$ through $\mathbb{E}_i[U_iV_iY_i]$ and is not a function of the marginals. The same construction shows that recording $\mathrm{Var}_i(\lambda_i)$ does not repair the state, since A and B may be re-weighted to match that moment too while (18) keeps σ_F^2 free. ■

Remark 3.5 (Recording $\mathrm{Var}_i(\lambda_i)$ Does Not Repair the State). $\mathrm{Var}_i(\lambda_i)$ is not a moment of $(\pi_i, y_i^{\mathrm{gap}})$, so it is not a coordinate of the state at issue, and matching it does not resolve the non-closure: A and B above can be matched on it too by a further, finer-grained construction, at the cost of one more free parameter each time. Each finite augmentation of the state is defeated by adding one more support point to the cross-section; no finite list of moments closes the recursion.

Lemma 3.12 (Multiplier Bound)**.** *Let* $X_i \in [0,1]$, $\bar{X} \equiv \mathbb{E}_i[X_i]$, $u_{\max} \equiv \max_i |X_i - \bar{X}|$, $a \equiv \bar{X} + u_{\max}$. *For any* W_i,

$$\mathrm{sd}_i(X_iW_i) \le a\,\mathrm{sd}_i(W_i) + u_{\max}|\bar{W}|.$$

Proof. $X_iW_i = \bar{X}W_i + U_iW_i$ with $U_i \equiv X_i - \bar{X}$; standard deviation is subadditive, giving $\mathrm{sd}_i(X_iW_i) \le \bar{X}\,\mathrm{sd}_i(W_i) + \mathrm{sd}_i(U_iW_i)$. Since $\mathrm{Var}_i(U_iW_i) \le \mathbb{E}_i[U_i^2W_i^2] \le u_{\max}^2\mathbb{E}_i[W_i^2] = u_{\max}^2(\mathrm{Var}_i(W_i)+\bar{W}^2)$ and $\sqrt{x+y} \le \sqrt{x}+\sqrt{y}$ for $x, y \ge 0$, $\mathrm{sd}_i(U_iW_i) \le u_{\max}(\mathrm{sd}_i(W_i) + |\bar{W}|)$. Combining gives the claim. ■

Proposition 3.3 ($\underline{F}_3$: A Lower Bound on the Coordination Threshold, and Certified Boundedness Below It)**.** *Theorem 3.1 blocks an exact recursion on the cross-sectional variances; the obstruction is equalities, not inequalities, so the coordination threshold can still be bounded from below. Let $n_\pi(t) \equiv \sqrt{V_\pi(t)}$ and let $a(F)$ be the scalar gain in the upper recursion*

$$n_\pi(t) \;\le\; a_\pi a(F)\, n_\pi(t-1) + a_\pi\big[F + u_{\max}(F)\pi^*\big] \tag{20}$$

delivered by Lemma 3.12 under Assumption 3.10, with $a(F) = \bar{X}(F) + u_{\max}(F) \ge \bar{X}(F)$ pointwise and $u_{\max}(F) < \infty$ on $[0, F_1)$. Define

$$\underline{F}_3 \;\equiv\; \inf\big\{F \in [0, F_1) :\; a_\pi a(F) \ge 1\big\}, \qquad \inf\emptyset \equiv F_1. \tag{21}$$

(i) Ordering — benchmark. *Under the sign benchmark* (P-odd)+(D-i)+(D-ii)*, and provided the coordination threshold F_3^* of Proposition 3.4 satisfies $F_3^* < F_1$: $\underline{F}_3 \le F_3^*$.*

(ii) Certified inflation boundedness — sign-free. *Under Assumptions 3.10, 3.2, 3.1 and 3.9 — the last supplying the bounded centering $|\bar{Y}| = |\bar{\pi}(t-1) - \pi^*| = \pi^*$ through Lemma 3.10, an information restriction and not a sign restriction: for every $F < \underline{F}_3$, $\sup_t V_\pi(t) < \infty$.*

(iii) Certified real boundedness — sign-free. *Under the hypotheses of (ii) alone: for every $F < \underline{F}_3$, $\sup_t \mathrm{Var}_i(y_{it}^{\mathrm{gap}}) < \infty$.*

Proof. (i) At F_3^* the crossing condition $\mu_c(F_3^*) = 1$ reads $a_\pi\bar{X}(F_3^*) = 1$. Since $a(F) \ge$

$\bar{X}(F)$ pointwise, $a_\pi a(F_3^*) \geq 1$, so F_3^* belongs to the set whose infimum defines $\underline{F}_3$ in (20); an infimum lies weakly below every element of its set, giving $\underline{F}_3 \leq F_3^*$. No monotonicity of a, $\bar{X}$ or $u_{\max}$ is used and no inverse function is formed.

(ii) For $F < \underline{F}_3$ the defining infimum gives $a_\pi a(F) < 1$. Assumption 3.1 makes $\mathbb{E}_{i,t-1}[R_t]$ cross-sectionally undispersed, so the injection term in (19) is finite on $[0, F_1)$. A scalar affine recursion with contraction factor strictly below one and bounded injection has bounded orbit: $\sup_t n_\pi(t) \leq a_\pi[F + u_{\max}(F)\pi^*]/(1 - a_\pi a(F)) + n_\pi(0) < \infty$, hence $\sup_t V_\pi(t) < \infty$. Only the upper recursion enters; no lower recursion and no covariance sign is used, and an upper recursion is never used here to claim divergence.

(iii) Lemma 3.1's IS block gives $\mathrm{Var}_i(y_{it}^{\mathrm{gap}}) \leq 2\sigma_0^2\big[F^2 + \sigma_F^2(t)\big]$, and Lemma 3.12 gives $\sigma_F^2(t) = \mathrm{Var}_i(X_i Y_i) \leq a(F)^2\big[V_\pi(t-1) + \bar{Y}^2\big]$. Both right-hand sides are bounded uniformly in t by (ii) and by Assumption 3.1: boundedness transfers from inflation to the output gap through the corrected IS equation, not by assumption. ■

Remark 3.6 (What $\underline{F}_3$ Does and Does Not Certify). $\underline{F}_3$ is a conservative certificate: strictly below it, inflation and output-gap dispersion are provably bounded. It is not a ceiling above which coordination fails, and this proposition classifies nothing above $\underline{F}_3$. Bounded dispersion is not exact coordination: Theorem 4.1 denies CE(0) at every $F > 0$, including on $[0, \underline{F}_3)$, so no sentence may read $F < F_3^* \Rightarrow$ CE. Four objects remain distinct: the universal exact-CE impossibility, the information-capacity ceiling F_1, the benchmark expectation-dispersion threshold F_3^*, and the bounded-dispersion certificate $\underline{F}_3$. The superseded label F_3^{ub}, which named this same object an *upper* bound, is retired: $\underline{F}_3$ is a lower bound on the coordination threshold and a boundedness certificate below itself, never a ceiling.

Lemma 3.13 (The Credibility-Weighted Dispersion Floor)**.** *Under MLRP*

(Proposition 2.1(ii)) and the symmetric cross-section of D7, for every $k \geq 1$,

$$\left|\frac{\mathbb{E}_i[X_i^k \theta_i^2]}{\mathbb{E}_i[X_i^k \theta_i]}\right| \geq \frac{F}{\sqrt{2}}.$$

At $k = 0$ the ratio is not defined: $X_i^0 \equiv 1$, so $\mathbb{E}_i[X_i^0 \theta_i] = \mathbb{E}_i[\theta_i] = 0$ by symmetry alone, independently of credibility.

Proof. Write $A_k(\theta) \equiv X(\theta)^k + X(-\theta)^k$, $D_k(\theta) \equiv X(-\theta)^k - X(\theta)^k$. By symmetry, $\mathbb{E}_i[X_i^k \theta_i^2] = \int_0^\infty \theta^2 A_k \, d\mu^+$ and $-\mathbb{E}_i[X_i^k \theta_i] = \int_0^\infty \theta D_k \, d\mu^+$, where μ^+ is μ restricted to $[0, \infty)$. MLRP gives $D_k \geq 0$ (strictly positive on a set of positive measure for $k > 0$), so the denominator does not change sign. Since $0 \leq D_k \leq A_k$, Cauchy–Schwarz in the measure $A_k \, d\mu^+$ gives $\int \theta D_k \, d\mu^+ \leq \left(\int \theta^2 A_k \, d\mu^+\right)^{1/2} \left(\int A_k \, d\mu^+\right)^{1/2}$, so the ratio in question is at least $\sqrt{m_2(k)}$, $m_2(k) \equiv \int \theta^2 A_k \, d\mu^+ / \int A_k \, d\mu^+$. Bounding $X(-\theta)^k \leq A_k \leq 2X(-\theta)^k$ on $[0, \infty)$ (valid since $X(\theta) \leq X(-\theta)$ there) gives $m_2(k) \geq \frac{1}{2} \int \theta^2 w_k \, d\mu^+ / \int w_k \, d\mu^+$, $w_k \equiv X(-\theta)^k$. Both θ^2 and w_k are non-decreasing on $[0, \infty)$ (the latter because X is decreasing), so Chebyshev's association inequality on the probability measure $\mu^+ / \mu^+([0, \infty))$ gives $\int \theta^2 w_k \, d\mu^+ / \int w_k \, d\mu^+ \geq \mathbb{E}_i[\theta_i^2] = F^2$, using symmetry for the last equality. Hence $m_2(k) \geq F^2/2$ for every k, giving the stated bound. ■

Remark 3.7 (Scope of the Dispersion Floor). Lemma 3.13 is exact and unconditional; it does not, by itself, extend the comonotone-propagation route of Lemma 3.10 to an exact closed-form threshold, since doing so additionally requires bounding the realized cross-sectional average $\langle X_j Y_j(t-1) \rangle$ uniformly over t, which depends on initial conditions not restricted by Proposition 3.3. Proposition 3.3's $\underline{F}_3$ remains the rigorous, unconditional result for this part of the Coordination Threshold.

Remark 3.8 (Low-Persistence Permanent Floor). For $\mu_c(F) < 1$, the recursion

satisfies $\liminf_{t\to\infty} \sigma_F^2(t) \geq \underline{I}(F)/(1-\mu_c(F)) > 0$: a strictly positive permanent floor binds even when fragmentation persistence is insufficient to generate explosive divergence. Theorem 4.1's Coordination Threshold result applies strictly to the A-RW persistent-fragmentation regime where $\mu_c(F) > 1$; Proposition 3.3's $\underline{F}_3$ does not depend on A-RW.

Assumption 3.11 (Credibility Attenuation — Benchmark Domain)**.** On $[0, F_1)$, the projected mean credibility $\bar{\lambda}^-(F) \equiv \int_0^1 \lambda_i^- \, di$ is continuous and non-increasing in F, with strict decrease on a dense subset; equivalently $\bar{X}(F) \equiv 1 - \bar{\lambda}^-(F)$ is continuous and non-decreasing, strictly increasing on a dense subset. This states directly, as an explicit benchmark restriction, the monotonicity that a stationary Bayesian updating rule with credibility loss increasing in signal noise would deliver; it is not derived from MLRP alone, which gives only the level result $g(F,t) > 0$ of Lemma 3.7, not monotonicity of the mean.

Proposition 3.4 (The Coordination Threshold as a First Crossing)**.** *Under Assumption 3.11, with $\mu_c(F) = a_\pi^2 \bar{X}(F)^2$ as in* (13)*, define the coordination threshold as the first crossing*

$$F_3^* \equiv \inf\big\{F \in [0, F_1) : \ \mu_c(F) \geq 1\big\}, \qquad \inf \emptyset \equiv F_1. \tag{22}$$

(i) Necessity. *$F_3^* < F_1$ requires $\kappa_0 \sigma_0 > 1 - \beta$, since $\bar{X}(F) \leq 1$ bounds $\mu_c(F) \leq a_\pi^2$. The condition is necessary and not sufficient.* (ii) Existence and interiority. *If in addition the endpoint condition*

$$\mu_c(0) < 1 < \mu_c(F_1^-) \tag{23}$$

holds, then $F_3^* \in (0, F_1)$*, continuity of* $\bar{X}$ *gives the crossing equation*

$$\bar{X}(F_3^*) = \frac{1-\beta}{\kappa_0 \sigma_0}, \tag{24}$$

and $\mu_c(F) < 1$ *for every* $F < F_3^*$. (iii) Divergence above the crossing. *At every* F *with* $\mu_c(F) > 1$*, under Assumption 3.2 and the sign layer of Lemma 3.9, iterating the lower recursion* $\sigma_F^2(t) \geq \mu_c(F)\sigma_F^2(t-1) + \mu_c(F)F^2$ *gives* $\sigma_F^2(t) \to \infty$*, so CE*(ε) *fails for every* $\varepsilon > 0$*. Divergence is claimed only from the lower recursion.* (iv) Uniqueness is not claimed. *Assumption 3.11 delivers* $\mu_c(F) \geq 1$ *on* $[F_3^*, F_1)$*; absent strict monotonicity throughout, the solution set of* $\mu_c(F) = 1$ *may be a non-degenerate interval and only its left endpoint is characterized. Every downstream statement uses the first crossing.* (v) Ordering. *Under* (22) *and Lemma 3.5,* $0 < \underline{F}_3 \leq F_3^* < F_1 < \infty$.

Proof. On $[0, F_1)$, Lemma 3.6 holds $\kappa_0\sigma_0$ constant, so $\mu_c = a_\pi^2 \bar{X}(F)^2$ inherits continuity and non-decrease from $\bar{X}$ (Assumption 3.11). (i) follows from $\bar{X} \leq 1$. Under (22) the set in (21) is non-empty and excludes a neighbourhood of 0, so its infimum is interior; continuity gives $\mu_c(F_3^*) = 1$, equivalently (23) since $\kappa_0\sigma_0 > 0$ and $\bar{X} > 0$; non-decrease gives $\mu_c < 1$ strictly to the left. (iii) is iteration of a lower recursion with gain above one. (iv) records that no step used strict monotonicity. (v) combines Proposition 3.3(i) with $F_1 < \infty$. ■

Corollary 3.3 (Rational Expectations: Exact Solution and the Absence of an Interior Threshold)**.** *Impose the rational-expectations benchmark* $\mathbb{E}_{i,t-1}[\pi_{i,t+1}] = \pi_i(t-1)$ *inside the forward-gap benchmark of Assumption 3.2 (*$h_i^+ = 0$*,* f_i *unrestricted). Assumption 3.3 is* not *imposed and cannot be: rational expectations force* $f_i = \pi_i(t-1) - \bar{\pi}(t-1)$*, which is not identically zero when* $F > 0$.

(a) The credibility weight drops out without a pointwise claim. *Recursive anchoring (Assumption 3.10) re-timed to filtration* $t-1$ *gives* $\mathbb{E}_{i,t-1}[\pi_{i,t+1}] - \pi^* = X_i Y_i$

with $X_i = 1-\lambda_i^-$ *and* $Y_i = \pi_i(t-1)-\pi^*$. *Rational expectations give* $X_iY_i = Y_i$, *hence* $\lambda_i^- Y_i = 0$ *region by region. The pointwise conclusion* $\lambda_i^- = 0$ *does* not *follow and is not needed:* $X_iY_i = (1-\lambda_i^-)Y_i = Y_i - \lambda_i^- Y_i = Y_i$ *identically, including on* $\{Y_i = 0\}$, *so*

$$\sigma_F^2(t) = \mathrm{Var}_i(X_iY_i) = \mathrm{Var}_i(Y_i) = V_\pi(t-1) \quad \textit{exactly}, \tag{25}$$

with no measure-theoretic side condition, no normalization $\bar{X} \equiv 1$, *and no appeal to Assumption 3.4 or to any covariance sign.*

(b) The cross-section solves in closed form. *Writing* $P_{it} \equiv \pi_{it} - \bar{\pi}_t$, *Lemma 3.1 with* $h_i^+ = 0$ *and* $f_i = P_{i,t-1}$ *delivers the exact difference equation*

$$P_{it} = a_\pi\big(q_i + P_{i,t-1}\big), \qquad a_\pi = \frac{\kappa_0\sigma_0}{1-\beta}, \tag{26}$$

an identity, not an inequality, whose unique solution from any initial cross-section $P_{i,-1}$ *is*

$$P_{it} = a_\pi^{t+1}P_{i,-1} + q_i\, a_\pi\, \frac{1-a_\pi^{t+1}}{1-a_\pi} \quad (a_\pi \neq 1), \qquad P_{it} = P_{i,-1} + (t+1)q_i \quad (a_\pi = 1). \tag{27}$$

(c) Three exhaustive cases in $\Delta \equiv \kappa_0\sigma_0 - (1-\beta)$. *Let the initial cross-section satisfy* $\mathrm{Cov}_i(q_i, P_{i,-1}) \geq 0$, *which holds trivially at an initially homogeneous cross-section* $P_{i,-1} \equiv 0$. *Then:*

$$\Delta < 0 \ (a_\pi < 1): \qquad V_\pi(t) \longrightarrow \Big(\frac{a_\pi}{1-a_\pi}\Big)^2 F^2 < \infty, \qquad \sigma_F^2(t) = V_\pi(t-1) \ \textit{bounded};$$

$$\Delta > 0 \ (a_\pi > 1): \qquad V_\pi(t) \to \infty \ \textit{at exact geometric rate} \ a_\pi^2 \ \textit{for every} \ F > 0;$$

$$\Delta = 0 \ (a_\pi = 1): \qquad V_\pi(t) = \mathrm{Var}_i\big(P_{i,-1} + (t+1)q_i\big) \to \infty \ \textit{for every} \ F > 0.$$

Boundedness in the first case is read off the exact solution (26)*, never from a lower recursion with gain below one; divergence in the second and third is read off the same exact solution, never from an upper recursion.*

(d) No interior threshold. *The effective feedback coefficient is* $\mu_c^{\mathrm{RE}} = a_\pi^2$*, constant in* F*. The equation* $\mu_c(F) = 1$ *selects no fragmentation level, so under rational expectations no interior coordination threshold* F_3^* *exists. The statement* $F_3^*(\mathrm{RE}) = 0$ *is false and is withdrawn: under* $\Delta > 0$ *divergence occurs at every* $F > 0$ *because the rate is flat in* F*, not because a threshold has migrated to the origin. The knife-edge lies in* $(\kappa_0, \sigma_0, \beta)$*, not in* F*. The universal finite-horizon floor* $A(F,t) > 0$ *of Lemma 3.9 holds at every* $F > 0$ *and every finite* t *regardless of* Δ*, so CE*(ε) *fails for every* $\varepsilon < A(F,t)$ *in all three cases.*

Proof. (a) is the displayed algebra. (b) Under Assumption 3.2, Lemma 3.1 reads $(1-\beta)P_{it} = \kappa_0\sigma_0(q_i + f_i)$; rational expectations set $f_i = \mathbb{E}_{i,t-1}[\pi_{i,t+1}] - \int_0^1 \mathbb{E}_{j,t-1}[\pi_{j,t+1}]\,dj = \pi_i(t-1) - \bar{\pi}(t-1) = P_{i,t-1}$, giving (25); (26) is its unique solution by direct iteration. (c) Apply Var_i to (26) with $\mathrm{Var}_i(q_i) = F^2$: the cross term carries the coefficient $2a_\pi^{t+1}a_\pi(1 - a_\pi^{t+1})/(1 - a_\pi) \geq 0$ under $\mathrm{Cov}_i(q_i, P_{i,-1}) \geq 0$, and for $a_\pi < 1$ both the initial and cross terms vanish geometrically, leaving the stated limit; for $a_\pi > 1$ the q_i coefficient satisfies $|a_\pi(1 - a_\pi^{t+1})/(1 - a_\pi)| \to \infty$ at rate a_π^{t+1}, and for $a_\pi = 1$ it equals $t + 1$; with $F > 0$ each diverges. (d) is immediate from (24) and (25). ■

Remark 3.9 (Reading the Rational-Expectations Branch). Three claims must stay separate. Universal reach failure (Theorem 4.1) holds at every $F > 0$ for every admissible mechanism and uses no expectations protocol. The benchmark coordination threshold F_3^* exists only when credibility attenuates with fragmentation (Assumption 3.11) and disappears under rational expectations, where the rate is flat. Divergence under rational expectations is a statement about $(\kappa_0, \sigma_0, \beta)$ alone. In

particular, $\Delta < 0$ delivers bounded dispersion but not exact coordination: $A(F,t) > 0$ still binds.

Proposition 3.5 (Threshold Ordering)**.** *Under the endpoint condition* (22)*,* $0 < F_3^* < F_1 < \infty$*.*

Proof. Proposition 3.4(ii) establishes $F_3^* \in (0, F_1)$ under (22) (which entails the necessary condition $\kappa_0\sigma_0 > 1-\beta$ of part (i)); Lemma 3.5 establishes $F_1 < \infty$ under the non-degeneracy condition. The strict ordering $0 < F_3^* < F_1 < \infty$ follows immediately. ■

4. The Monetary Reach Theorem

Under MTP-3, $\mathrm{Var}_i(\mathbb{E}_{i,t-1}[\pi_{it}]) = \mathrm{Var}_i(\pi_{it})$ identically — the same idempotence used throughout Section 3 — so this object cannot serve as an independent expectation state on the right-hand side of a decomposition of $\mathrm{Var}_i(\pi_{it})$ itself. Write

$$V_\pi(t) \equiv \mathrm{Var}_i(\pi_{it}), \qquad \sigma_F^2(t) \equiv \mathrm{Var}_i\big(\mathbb{E}_{i,t-1}[\pi_{i,t+1}]\big),$$

the first realized inflation dispersion, the second one-step-ahead forecast dispersion measured at $t-1$; the symbol $\sigma_E^2(t)$ is retired.

Proposition 4.1 (Inflation Dispersion Identity)**.** *Under Lemma 3.1, squaring* (6) *and integrating over regions gives, with* $q_i \equiv r_i^* - r^*$*,*

$$V_\pi(t) = \frac{\kappa_0^2}{(1-\beta)^2}\Big[\sigma_0^2 F^2 + \mathrm{Var}_i(h_i^+) + \sigma_0^2\,\mathrm{Var}_i(f_i) + 2\sigma_0\,\mathrm{Cov}_i(q_i, h_i^+) + 2\sigma_0^2\,\mathrm{Cov}_i(q_i, f_i) + 2\sigma_0\,\mathrm{Cov}_i(h_i^+, f_i)\Big]. \quad (28)$$

Under Assumption 3.2 ($h_i^+ = 0$), this reduces to

$$V_\pi(t) = a_\pi^2\Big[F^2 + \mathrm{Var}_i(f_i) + 2\,\mathrm{Cov}_i(q_i, f_i)\Big] + \Sigma(F), \qquad a_\pi \equiv \frac{\kappa_0\sigma_0}{1-\beta}, \tag{29}$$

where $\Sigma(F) \equiv \mathrm{Var}_i(\varepsilon_{it}) = \sigma_\varepsilon^2 \geq 0$ is the additive idiosyncratic component of realized inflation carried by (D5); the first bracket is the exact solution of (6) *and Σ is added, not assumed away.*

Proof. $\mathbb{E}_i[q_i] = \mathbb{E}_i[h_i^+] = \mathbb{E}_i[f_i] = 0$ by construction; squaring (6) and integrating over i gives the six displayed terms directly, with no term dropped or assumed signed. ■

Lemma 4.1 (Reduced Cross-Section)**.** *Under Lemma 3.1, for every i and t,*

$$\pi_{it} - \bar\pi_t = a_\pi\big(q_i + w_i\big), \qquad w_i \equiv f_i + \frac{h_i^+}{\sigma_0}, \qquad a_\pi = \frac{\kappa_0\sigma_0}{1-\beta}, \tag{30}$$

for the systematic component of (D5), hence $V_\pi(t) = a_\pi^2\,\mathrm{Var}_i(q_i + w_i) + \sigma_\varepsilon^2$ exactly, the additive innovation raising every floor below and lowering none. Under Assumption 3.2, $w_i = f_i$.

Proof. Divide the identity $(1-\beta)(\pi_{it} - \bar\pi_t) = \kappa_0\sigma_0 q_i + \kappa_0 h_i^+ + \kappa_0\sigma_0 f_i$ of Lemma 3.1 by $1-\beta > 0$ and collect $\kappa_0\sigma_0$. No sign restriction and no expectations protocol is used: the step is division of an identity. ■

Proposition 4.2 (The PE–US Quantitative Floor)**.** *Let $F > 0$.* (i) Sharp condition. *The floor*

$$V_\pi(t) \;\geq\; a_\pi^2 F^2 \tag{31}$$

holds if and only if

$$2\,\mathrm{Cov}_i(q_i, w_i) + \mathrm{Var}_i(w_i) \;\geq\; 0. \tag{32}$$

The floor is therefore not universal, and (31) *is the weakest hypothesis that delivers it: the theorem is exactly as strong as its proof.* (ii) Benchmark sufficiency. *Under Assumption 3.2 ($w_i = f_i$), the covariance sign* $\mathrm{Cov}_i(q_i, f_i) \geq 0$ *— established by Lemma 4.2 under* (P-odd)+(D-i)+(D-ii)*, and imposable directly under Assumption 3.4 where that symmetry route is unavailable — implies* (31) *and hence* (30). *Under Assumption 3.3 ($w_i = 0$) the floor holds with equality, $V_\pi(t) = a_\pi^2 F^2$.* (iii) Universal degraded floor. *With no sign hypothesis whatsoever, the L^2 triangle inequality applied to* (29) *gives, for every admissible mechanism,*

$$\sqrt{V_\pi(t)} \;\geq\; a_\pi\big(F - \mathrm{sd}_i(w_i)\big), \tag{33}$$

so $V_\pi(t) \geq a_\pi^2\big(F - \mathrm{sd}_i(w_i)\big)^2 > 0$ whenever forecast-deviation dispersion falls short of natural-rate dispersion, $\mathrm{sd}_i(w_i) < F$. Driving the floor to zero requires exact offsetting, $w_i = -q_i$ almost everywhere: the regional forecast field must reproduce the entire unspanned state, which (D1)–(D3) deny to any scalar instrument. (iv) PE–US separation. *Whenever* (31) *holds, US(η) fails for every $\eta < a_\pi^2 F^2$ at every $F \in (0, F_1)$, while PE holds at the constant transmission floor $\kappa(F)\sigma(F) = \kappa_0\sigma_0 > 0$ (Lemma 3.6): transmission and uniform stabilization are separated by a closed-form margin proportional to the squared dynamic loading a_π^2, not to the static γ_A^2 and not to γ_I.*

Proof. (i) Expand $\mathrm{Var}_i(q_i + w_i) = F^2 + 2\,\mathrm{Cov}_i(q_i, w_i) + \mathrm{Var}_i(w_i)$ in (29); since $a_\pi^2 > 0$, (30) is equivalent to joint non-negativity of the last two terms, which is (31). Necessity and sufficiency are simultaneous because (29) is an identity, not a bound. (ii) $\mathrm{Cov}_i(q_i, f_i) \geq 0$ makes both terms in (31) non-negative; $w_i = 0$ makes them zero. (iii) $\mathrm{sd}_i(q_i + w_i) \geq \mathrm{sd}_i(q_i) - \mathrm{sd}_i(w_i)$ by the triangle inequality in L^2 of the spatial measure, and $\mathrm{sd}_i(q_i) = F$. (iv) Combine (30) with Lemma 3.6. ∎

Lemma 4.2 (Sign of the Expectation Covariance)**.** *Under Assumption 3.10,* $f_i = X_iY_i - \overline{X_jY_j}$ *with* $X_i \equiv 1-\lambda_i^-$, $Y_i \equiv \pi_i(t-1)-\pi^*$*; writing* $X_i = \bar{X}+U_i$, $Y_i = \bar{Y}+V_i$, $\mathrm{Cov}_i(q_i, f_i) = \bar{X}\,\mathrm{Cov}_i(q_i, V_i) + \bar{Y}\,\mathrm{Cov}_i(q_i, U_i) + \mathrm{Cov}_i(q_i, U_iV_i)$*. Under* (D-i) $\bar{\pi}(t-1) \le \pi^*$ *(established exactly by Lemma 3.10) and* (D-ii) *Assumption 3.8 with Lemma 3.8,* $\mathrm{Cov}_i(q_i, f_i) \ge 0$*, strictly whenever* $\mathrm{Var}_i(V_i) > 0$*. Absent (D-i) or (D-ii) the sign is not determined by comonotonicity:* X_i *is non-increasing in* q_i *while* Y_i *is increasing in* q_i*, so* X_iY_i *is not a comonotone function of* q_i *and Chebyshev's sum inequality does not apply directly.*

Proof. By Lemma 3.1, Y_i is increasing in q_i, so $\mathrm{Cov}_i(q_i, V_i) \ge 0$; the first term is non-negative since $\bar{X} > 0$. By Proposition 2.1(ii), $\Lambda' \ge 0$ makes X_i non-increasing in q_i, so $\mathrm{Cov}_i(q_i, U_i) \le 0$; under (D-i), $\bar{Y} \le 0$, making the second term non-negative. Under (D-ii), U_i and V_i are both odd in q_i (Lemma 3.8), so $q_iU_iV_i$ is odd and the third term vanishes. ■

Lemma 4.3 (Environmental Separation)**.** *The capacity metric* $h_{\mathrm{agg}}(F) = \int_0^1 h_i\,di$ *is bounded by the constraints of Strict Capacity Adherence condition (A) inside Definition 2.2. Every mechanism in* $\mathfrak{M}(\mathfrak{D})$ *treats the transmission triple* $(\kappa(F), \sigma(F), h_{\mathrm{agg}}(F))$ *as exogenous environmental parameters; altering them exits* $\mathfrak{M}(\mathfrak{D})$ *by violating Condition (C) of Definition 2.2. Augmenting the instrument with autoregressive lags* $\{R_{t-k}\}_{k\ge1}$ *does not expand spatial coverage: the structural state* $\theta \in \mathbb{R}^K$ *is time-invariant, so each lag* $R_{t-k} = \varphi(\Omega_{CB,t-k})$ *responds to the identical state and adds no independent orthogonal dimension; the kernel dimension* $K-1 \ge 1$ *is invariant to lag order under bounded capacity (P1).*

Proof. Conditions (A) and (C) of Definition 2.2 deliver the first two claims directly. Each lag $R_{t-k} \in \mathbb{R}$ is a scalar satisfying Conditions (B) and (D1)–(D3); the scalar codomain is unchanged, so $\dim(\ker \delta_{t-k}) = K-1 \ge 1$ at every lag order by the

rank-nullity theorem applied to Condition (D3). ■

Remark 4.1 (Structural Parametric Invariance). The deep parameters (κ_0, σ_0) — pricing rigidities and spatial elasticities — are properties of the physical environment, invariant to expectation-formation architecture (Lucas, 1976); the transmission bounds are invariant across $\mathfrak{M}(\mathfrak{D})$.

Scalar aggregation limits are dimensional, not operational (Arrow, 1951; Mundell, 1961). The classical Mundell-Fleming trilemma is the low-fragmentation linearization of the impossibility (Theorem 4.2): the classical trade-off is the degenerate limit of a dimensional collapse, not an independent constraint. No institutional assumption is required.

Theorem 4.1 (The Monetary Reach Theorem)**.** *No scalar-instrument monetary authority can simultaneously coordinate inflation expectations, transmit policy to the real economy, and stabilize outcomes uniformly across regions of a heterogeneous jurisdiction. Let $F > 0$. Under physical primitives* (P1) $\bar{H} < \infty$, (P2) $\mathrm{Var}_i(r_i^*) > 0$, *and* (P3) $u''(\bar{c}, \bar{n}) < 0$, *with the scalar-instrument dimensional constraints* (D1)–(D3) *and* (D5)–(D7)*, the finite-horizon CE-failure floor $A(F,t) > 0$ holds unconditionally at every finite t; the closed-form thresholds require additionally Assumptions 3.7, 2.3, 3.11. The bilateral CE, PE, US failure structure is as follows.*

Universal Infeasibility. *$\{CE, PE, US\}$ is jointly infeasible across every admissible mechanism in $\mathfrak{M}(\mathfrak{D})$. This exclusion operates through three logically orthogonal boundaries. (a) alone secures the headline: $A(F,t) > 0$ makes $CE(0)$ unattainable for every $F > 0$ and every mechanism, and $CE(0)$ failing suffices for $\{CE, PE, US\}$ jointly infeasible; (b) and (c) sharpen this floor using the canonical benchmark's transmission coefficients (Remark 3.4).*

(a) CE–PE bilateral. *Rank-nullity leaves $\gamma_u^2 \mathrm{Var}_i(r_i^*) > 0$ permanently unspanned*

for every admissible $\gamma_u > 0$, establishing the strict floor $A(F,t) > 0$ at every finite t. CE(ε) fails for every tolerance $\varepsilon < A(F,t)$, not merely at exact zero.

(b) PE–US bilateral. *Unspanned divergence floors spatial inflation variance at $a_\pi^2 \operatorname{Var}_i(r_i^*) > 0$ under Assumption 3.2 (Proposition 4.2), independently of instrument configuration. The loading is the dynamic a_π, never the static product $\kappa_0\sigma_0$: predetermination multiplies the floor by $(1-\beta)^{-2}$.*

(c) CE–US bilateral. *Any cross-sectional expectation dispersion $\sigma_F^2(t) > 0$ maps through the Phillips curve identity into a strict physical stabilization floor, independently of instrument scarcity:*

$$\eta \geq \operatorname{Var}_i(\pi_i - \pi^*) \geq a_\pi^2 F^2 > 0 \qquad \textit{whenever } 2\operatorname{Cov}_i(q_i, w_i) + \operatorname{Var}_i(w_i) \geq 0,$$

and, with no sign restriction whatsoever,

$$\sqrt{\operatorname{Var}_i(\pi_i - \pi^*)} \;\geq\; a_\pi\big[F - \operatorname{sd}_i(w_i)\big]$$

by the triangle inequality (Proposition 4.2), where $w_i \equiv f_i + h_i^+/\sigma_0$. US($\eta$) therefore fails for every η below the applicable floor. Jensen's inequality operating over the strictly concave regional utility (Primitive P3) forces aggregate welfare strictly below the centralized evaluation; the resultant welfare gap $W_J < 0$ scales proportionate to $\sigma_F^2(t)$, independently of the stabilization floors in (a) and (b).

The Coordination Threshold. *Restricting to the recursive-feedback subclass $\mathfrak{M}_{\mathrm{RF}}(\mathfrak{D})$ of Assumption 3.10 is necessary, not merely convenient (Proposition 3.2, Remark 3.3): no threshold binds coordination dispersion across the whole of $\mathfrak{M}(\mathfrak{D})$. Within $\mathfrak{M}_{\mathrm{RF}}(\mathfrak{D})$, given the two most favorable initial conditions $\bar\pi(-1) \leq \pi^*, \bar\pi(0) \leq \pi^*$ (Lemma 3.10), coordination survives below a rigorous, computable bound: for*

$F < \underline{F}_3$ (Proposition 3.3), both cross-sectional inflation dispersion and output-gap dispersion remain bounded at every date. Above F_1, CE fails via $A(F,t) > 0$ independently; the capacity constraint releases and aggregate attention decays continuously toward zero; for any transmission tolerance $\delta > 0$, no mechanism in $\mathfrak{M}(\mathfrak{D})$ sustains $\kappa(F)\sigma(F) \geq \delta$ above the finite threshold $F(\delta)$ (Corollary 3.2).

F_1 is the aggregate transmission ceiling (Lemma 3.5): an unconditional outer-bound independent of any coordination mechanism. The exact coordination threshold $F_3^* \in (0, F_1)$, closed-form in (41), holds conditional on the reconciliation of Remark 3.4; $\underline{F}_3 \leq F_3^*$ is unconditional and does not depend on that reconciliation.

The geometric feasibility set $\mathcal{F}(F) \equiv \{(CE, PE, US)$ simultaneously achievable$\}$ contracts as follows. For $F < \underline{F}_3$, CE(ε)+PE is subject only to the finite-horizon floor $A(F,t) > 0$; coordination dispersion is proven bounded (Proposition 3.3), though exact-zero dispersion is not achievable. Conditional on Remark 3.4's reconciliation, CE(ε)+PE loses viability above F_3^* and PE+US achieves strictly positive bounded floors throughout $F \in (F_3^*, F_1)$; between $\underline{F}_3$ and F_3^* this characterization is pending, not asserted. The entire operational manifold collapses above F_1. No monetary intervention restores viability above these coordinates; the constraints are dimensional, not behavioral, and neither institutional redesign nor rule modification moves the thresholds. The three-regime partition of the domain is illustrated in Figure 1, drawn under the conditional closed form of Remark 3.4.

A single scalar interest rate cannot simultaneously coordinate regional inflation expectations, transmit policy to real activity, and stabilize inflation uniformly across regions; persistent spatial fragmentation severs the central bank from uniform macroeconomic control regardless of institutional design.

Remark 4.2 (Universal and Benchmark-Specific Failure Modes). Theorem 4.1

separates two logically distinct mechanisms.

First (Universal Infeasibility): the impossibility of exact joint fulfillment is universal across every admissible mechanism in $\mathfrak{M}(\mathfrak{D})$ through the finite-horizon bound $A(F,t) > 0$ and the sign-free floor $\sqrt{\mathrm{Var}_i(\pi_i - \pi^*)} \geq a_\pi[F - \mathrm{sd}_i(w_i)]$ of Proposition 4.2(iii); the closed-form floor $\mathrm{Var}_i(\pi_i - \pi^*) \geq a_\pi^2 \mathrm{Var}_i(r_i^*)$ is benchmark-specific, holding exactly under condition (31). Each rests on bounded capacity (P1), irreducible heterogeneity (P2), and the scalar instrument's rank-deficiency (D1–D3) jointly, with positive transmission (MTP-1, MTP-2) additionally load-bearing for the inflation floor; no single primitive delivers either result alone.

Second (the Coordination Threshold): the explosive growth of expectation dispersion, $\sigma_F^2(t) \to \infty$, additionally requires the persistent-drift benchmark (Assumption 2.3) and the adaptive updating rule of Lemma 3.9. Persistent drift is what keeps the injection floor bounded away from zero as $t \to \infty$ rather than merely at each finite t; absent it, nothing here rules out $\bar{\xi}(t) \to \pi^*$ eventually, and $A(F,t)$ was only ever shown positive pointwise in t. The explosive result binds within the benchmark environment; the universal floor binds without it.

Remark 4.3 (The $F = 0$ Degenerate Boundary). Theorem 4.1 applies for $F > 0$. At $F = 0$, all three bilateral impossibilities collapse: $A(0,t) = 0$ because $g(0,0) = 0$ (Lemma 3.7); $\mathrm{Var}_i(\lambda_i) \to 0$ as $F \to 0^+$; and uniform stabilization is achievable under perfect homogeneity. The three objectives are jointly attainable at $F = 0$ and at $F = 0$ alone. This attainability is established for the general adaptive benchmark of Section 3; whether it survives the rational-expectations limit of Corollary 3.3, where $\mu_c^{\mathrm{RE}} = a_\pi^2$ is constant in F, is not verified here and is left open. For any $F > 0$, however small, all three bilateral impossibilities bind simultaneously.

Assumption 4.1 (Common Filtration at Homogeneity — Benchmark)**.** At $F = 0$ every region holds the same information set, $\Omega_{it} = \Omega_t$ for every i, and the aggregate natural rate is predetermined: $r_t^* = \rho + \zeta_t$ with $\zeta_t \in \Omega_{t-1}$, $\mathbb{E}[\zeta_t] = 0$, $\mathrm{Var}(\zeta_t) = s^2 > 0$. The central bank observes ζ_t only through $\hat{\zeta}_t = \zeta_t + e_t$, with e_t mean-zero, independent of ζ_t, and $\mathrm{Var}(e_t) = \sigma_e^2 > 0$; the strict positivity of σ_e^2 is the quantitative content of the bounded-capacity primitive (P1) and of Assumption 2.2. This is an information restriction, not a sign restriction, and it is not inferred from $q_i = 0$: vanishing natural-rate dispersion constrains natural rates, not filtrations. It is imposed only at $F = 0$ and is used nowhere in Theorem 4.1.

Theorem 4.2 (Exact Existence of the Classical Mundell–Fleming Frontier at $F = 0$)**.** *Let $F = 0$. Then $\mathrm{Var}_i(r_i^*) = 0$ with $\mathbb{E}_i[q_i] = 0$ forces $q_i = 0$ almost everywhere in i. Under Assumption 4.1 the cross-section collapses exactly: $h_i^+ = 0$ and $f_i = 0$ for every i, so Lemma 3.1 gives $\pi_{it} = \bar{\pi}_t$, and (3) with a common filtration and a common instrument gives $y_{it}^{\mathrm{gap}} = \bar{y}_t^{\mathrm{gap}}$ — equality across regions, not to zero, since $\sigma_e^2 > 0$ leaves the bank a strictly random aggregate forecast error (Lemma 2.2): the heterogeneous jurisdiction is a single representative economy, and a homogeneous equilibrium exists in the sense of Definition 2.1 by Lemma 2.1. Normalize $\pi^* = 0$ and take the scalar policy family $R_t = \rho + \pi^* + \phi\hat{\zeta}_t$, $\phi \in \mathbb{R}$, whose single coordinate ϕ is the exact policy coordinate at $F = 0$. Then, exactly,*

$$\pi_t = a_\pi(1-\phi)\zeta_t, \qquad y_t^{\mathrm{gap}} = \sigma_0\big[(1-\phi)\zeta_t - \phi e_t\big], \tag{34}$$

so the attainable pair $\big(V_\pi(\phi), V_y(\phi)\big) \equiv \big(\mathrm{Var}(\pi_t), \mathrm{Var}(y_t^{\mathrm{gap}})\big)$ is

$$V_\pi(\phi) = a_\pi^2(1-\phi)^2 s^2, \qquad V_y(\phi) = \sigma_0^2\big[(1-\phi)^2 s^2 + \phi^2\sigma_e^2\big]. \tag{35}$$

Eliminating ϕ on $[0,1]$ gives the frontier in closed form,

$$V_y = \frac{\sigma_0^2}{a_\pi^2}\, V_\pi + \sigma_0^2\sigma_e^2\Big(1 - \frac{\sqrt{V_\pi}}{a_\pi s}\Big)^2, \qquad V_\pi \in \big[0,\, a_\pi^2 s^2\big], \tag{36}$$

a set, not a tangent: it is attained exactly, it reaches $V_\pi = 0$ at the finite real cost $V_y = \sigma_0^2\sigma_e^2$, and its slope

$$\frac{dV_y}{dV_\pi} = \sigma_0^2\Big[\frac{1}{a_\pi^2} - \frac{\sigma_e^2\,\phi}{a_\pi s\sqrt{V_\pi}}\Big] \tag{37}$$

is strictly negative on a non-degenerate interval, tending to $-\infty$ as $\phi \to 1^-$ and equal to $\sigma_0^2/a_\pi^2 > 0$ at $\phi = 0$: the classical trade-off between price stability and real stabilization exists *at exactly $F = 0$, with strictly positive trade-off coefficients $\big(\sigma_0^2/a_\pi^2,\ \sigma_0^2\sigma_e^2\big)$. The Regime I loss* (4) *is minimized at the interior policy*

$$\phi^* = \frac{s^2\big(a_\pi^2 + \alpha\sigma_0^2\big)}{s^2\big(a_\pi^2 + \alpha\sigma_0^2\big) + \alpha\sigma_0^2\sigma_e^2} \in (0,1), \qquad L_{\mathrm{I}}(0) = \frac{\alpha\sigma_0^2\sigma_e^2\, s^2\big(a_\pi^2 + \alpha\sigma_0^2\big)}{s^2\big(a_\pi^2 + \alpha\sigma_0^2\big) + \alpha\sigma_0^2\sigma_e^2}, \tag{38}$$

so the frontier is attained at an interior point, not at a corner. It is non-degenerate if and only if $\sigma_e^2 > 0$: at $\sigma_e^2 = 0$, $\phi^* = 1$, $L_{\mathrm{I}}(0) = 0$, and the frontier collapses to the bliss point $(0,0)$. The classical trade-off is therefore the exact consequence of bounded capacity (P1) acting on a homogeneous state, and of nothing else.

Proof. $\mathrm{Var}_i(r_i^*) = 0$ and $\mathbb{E}_i[q_i] = 0$ give $q_i = 0$ a.e. Under Assumption 4.1 every conditional expectation is common across i, so h_i^+ and f_i, being differences between a region's forecast and the cross-sectional average of forecasts, vanish identically; this is a consequence of the stated filtration restriction, not of $q_i = 0$. Lemma 3.1 then gives $(1-\beta)(\pi_{it} - \bar\pi_t) = 0$, and $1 - \beta > 0$ gives $\pi_{it} = \bar\pi_t$. For the representative economy, MTP-3 makes π_t Ω_{t-1}-measurable, so (2) reads $(1-\beta)\pi_t = \kappa_0\mathbb{E}_{t-1}[y_t^{\mathrm{gap}}]$. Under the stated policy family, (3) gives $y_t^{\mathrm{gap}} = \mathbb{E}_t[y_{t+1}^{\mathrm{gap}}] + \sigma_0\big(r_t^* + \mathbb{E}_t[\pi_{t+1}] - R_t\big)$; the stationary forward terms are constants absorbed into the policy intercept, leaving

$y_t^{\text{gap}} = \sigma_0(\zeta_t - \phi\hat{\zeta}_t) = \sigma_0[(1-\phi)\zeta_t - \phi e_t]$. Since $\zeta_t \in \Omega_{t-1}$ and $\mathbb{E}_{t-1}[e_t] = 0$, $\mathbb{E}_{t-1}[y_t^{\text{gap}}] = \sigma_0(1-\phi)\zeta_t$ and $\pi_t = \frac{\kappa_0\sigma_0}{1-\beta}(1-\phi)\zeta_t = a_\pi(1-\phi)\zeta_t$, which is (33). Taking variances and using independence of ζ_t and e_t gives (34); $1-\phi = \sqrt{V_\pi}/(a_\pi s)$ on $\phi \leq 1$ gives (35) and, on differentiating, (36). Substituting (34) into (4) gives $L_{\text{I}}(\phi) = (1-\phi)^2 s^2\big(a_\pi^2 + \alpha\sigma_0^2\big) + \alpha\sigma_0^2\sigma_e^2\phi^2$, strictly convex, with first-order condition $-(1-\phi)s^2(a_\pi^2 + \alpha\sigma_0^2) + \alpha\sigma_0^2\sigma_e^2\phi = 0$ delivering (37), interior because both $s^2(a_\pi^2 + \alpha\sigma_0^2)$ and $\alpha\sigma_0^2\sigma_e^2$ are strictly positive; substitution gives the displayed minimized loss. ■

Proposition 4.3 (Destruction of the Classical Interior Frontier at Every $F > 0$)**.** *Let $F > 0$, retain the full-anchoring benchmark Assumption 3.3 ($h_i^+ = 0$ and $f_i = 0$, so $w_i = 0$), and retain the scalar policy family of Theorem 4.2. The stronger hypothesis is stated because the proof uses $f_i = 0$ and not merely $h_i^+ = 0$; the weaker forward-gap benchmark Assumption 3.2 alone does not deliver $\pi_{it} - \bar{\pi}_t = a_\pi q_i$. Then, exactly,*

$$V_\pi(\phi; F) = a_\pi^2\big[(1-\phi)^2 s^2 + F^2\big], \qquad V_y(\phi; F) = \sigma_0^2\big[(1-\phi)^2 s^2 + \phi^2\sigma_e^2 + F^2\big], \quad (39)$$

so the attainable set is the $F = 0$ frontier translated by the fixed vector $\big(a_\pi^2 F^2,\ \sigma_0^2 F^2\big)$, on which the instrument has no purchase: $\partial\big(a_\pi^2 F^2, \sigma_0^2 F^2\big)/\partial\phi = 0$. (i) Unspanned block. *The Jacobian of the attainable map $\phi \mapsto (V_\pi, V_y)$ has rank one, while the outcome space carries a component — the dispersion block — on which it is identically zero. This is the scalar-versus-multidimensional rank deficiency of (D1)–(D3) in exact quantitative form, and it is the same deficiency Theorem 4.1 invokes.* (ii) Non-existence of the classical frontier. *The defining property of the classical frontier — price stability attainable at finite real cost, $V_\pi = 0$ with $V_y < \infty$ — fails: $\inf_\phi V_\pi(\phi; F) = a_\pi^2 F^2 > 0$ for every scalar policy, and no reparametrization of ϕ restores it because the deficit is ϕ-invariant. There is no interior classical frontier at any $F > 0$; Proposition 4.2 records the same deficit as the PE–US floor.* (iii)

Knife-edge. *The map* $F \mapsto \inf_\phi V_\pi$ *equals* $a_\pi^2 F^2$*: zero at exactly* $F = 0$ *and strictly positive at every* $F > 0$*. The classical Mundell–Fleming trade-off exists at* $F = 0$ *and at* $F = 0$ *alone; for every* $F > 0$ *the classical interior frontier does not exist and Theorem 4.1's impossibility applies. The frontier is not destroyed at* $F = 0$*.* (iv) What survives without full anchoring. *Drop* $f_i = 0$ *and retain only Assumption 3.2. Then* $w_i = f_i$ *and Lemma 4.1 gives* $V_\pi(\phi; F) = a_\pi^2\big[(1-\phi)^2 s^2 + \mathrm{Var}_i(q_i + f_i)\big] + \sigma_\varepsilon^2$*, so the* ϕ*-invariant deficit becomes* $a_\pi^2 \mathrm{Var}_i(q_i + f_i) + \sigma_\varepsilon^2$*, which the sign-free bound of Proposition 4.2 bounds below by* $a_\pi^2\big[F - \mathrm{sd}_i(f_i)\big]^2 + \sigma_\varepsilon^2$*. Rank deficiency in (i) is unaffected, since* $\partial/\partial\phi$ *annihilates the deficit whatever its magnitude; only the exact value* $a_\pi^2 F^2$ *in (ii)–(iii) requires* $f_i = 0$*. The knife-edge at* $F = 0$ *is therefore benchmark-specific in its* magnitude *and universal in its* rank*.*

Proof. Under Assumption 3.3, Lemma 4.1 gives $\pi_{it} - \bar\pi_t = a_\pi q_i$ since $w_i = 0$, and (3) with R_t common across i gives $y_{it}^{\mathrm{gap}} - \bar y_t^{\mathrm{gap}} = \sigma_0 q_i$. Adding these cross-sectional components to the aggregate outcomes of Theorem 4.2, and using orthogonality of the cross-sectional deviation q_i to the aggregate innovations (ζ_t, e_t) under Assumption 2.1, gives (38) with $\mathrm{Var}_i(q_i) = F^2$. Claims (i)–(iii) read off the displayed expressions. For (iv), drop $f_i = 0$: Lemma 4.1 gives $\pi_{it} - \bar\pi_t = a_\pi(q_i + f_i)$ plus the (D5) innovation, whose cross-sectional variance is $a_\pi^2 \mathrm{Var}_i(q_i + f_i) + \sigma_\varepsilon^2$ and is independent of ϕ by the same orthogonality; the displayed lower bound is Proposition 4.2's universal degraded floor applied to $w_i = f_i$. ∎

Corollary 4.1 (Quantitative Bounds)**.** *Under the conditions of Theorem 4.1, for any* $F > 0$*: CE and PE together imply*

$$\mathrm{Var}_i(\pi_i - \pi^*) \geq a_\pi^2 \mathrm{Var}_i(r_i^*) + \sigma_\varepsilon^2 > 0, \qquad a_\pi = \frac{\kappa_0 \sigma_0}{1-\beta}, \tag{40}$$

under Assumption 3.2, which removes the forward-gap channel that exact CE does

not reach. For $F \in (0, F_1)$, US(η) fails for all $\eta < a_\pi^2 \operatorname{Var}_i(r_i^*)$. Above F_1, aggregate extraction leaves the capacity boundary; for any transmission tolerance $\delta > 0$, PE(δ) fails above the finite threshold $F(\delta)$ defined by $\kappa(F(\delta))\sigma(F(\delta)) = \delta$, with $\lim_{F\to\infty} \kappa(F)\sigma(F) = 0$. Uniform stabilization is unattainable independently of bounded parameter variances.

By the universal floor $A(F,t) > 0$ (Lemma 3.7), PE and US together violate CE through three layers: $A(F,t) = g(F,t)(\pi^* - \bar{\xi}(t))^2 > 0$ at every $F > 0$ and every finite t from the scalar-instrument restriction alone; under outcome-based mechanisms and A-RW, $\sigma_F^2(t)$ grows at rate $\mu_c(F) > 1$ above F_3^*; and the variance recursion operator $\mathcal{T}_\sigma : \sigma_F^2(t) \mapsto \sigma_F^2(t+1)$ is expansive for $F > F_3^*$, conditional on Remark 3.4 (Proposition 3.4); unconditionally, the same operator is bounded below $\underline{F}_3$ (Proposition 3.3).

Proof. Under exact CE ($\varepsilon = 0$), $\operatorname{Var}_i(\mathbb{E}_i[x]) = 0$ and $W_J = 0$: the Jensen welfare-deficit channel is inactive. Exact CE(0) requires $A(F,t) = 0$, which fails for all $F > 0$ and all finite t (Lemma 3.7: $g(F,t) = \operatorname{Var}_i(\lambda_i) > 0$). The CE–US bilateral impossibility is therefore established from the scalar-instrument restriction alone, without invoking W_J; the Jensen channel ($W_J \neq 0$) provides the independent SEC-3 mechanism operating under any imperfect CE.

Instrument scarcity. Simultaneous regional stabilization requires $R_t = \pi^* + r_i^*$ for every $i \in [0,1]$; with $\operatorname{Var}_i(r_i^*) > 0$ (P2), the target constitutes a non-constant element of $L^2([0,1])$. The scalar-instrument restriction (D1)–(D3) confines the instrument to $R_t \in \mathbb{R}$; the map $\mathbb{R} \to L^2([0,1])$ defined by $R_t \mapsto R_t \cdot \mathbf{1}_{[0,1]}$ has range equal to the one-dimensional constant-function subspace of $L^2([0,1])$. The target $\{\pi^* + r_i^*\}_{i\in[0,1]}$ lies outside this subspace whenever $\operatorname{Var}_i(r_i^*) > 0$. The residual from the best scalar

approximation satisfies

$$\inf_{R_t \in \mathbb{R}} \mathrm{Var}_i(r_i^* - (R_t - \pi^*)) = \mathrm{Var}_i(r_i^*) > 0,$$

which the Phillips block of Lemma 3.1 transmits at the dynamic loading, delivering the variance floor $\mathrm{Var}_i(\pi_i - \pi^*) \geq a_\pi^2 \, \mathrm{Var}_i(r_i^*) > 0$ under Assumption 3.2.

Quantitative floor. Exact CE sets $\mathrm{Var}_i(f_i) = 0$, and Assumption 3.2 sets $h_i^+ = 0$; identity (27) then reduces exactly to $\mathrm{Var}_i(\pi_i - \pi^*) = a_\pi^2 \, \mathrm{Var}_i(r_i^*) + \Sigma(F)$, which is (39). The loading is the dynamic $a_\pi = \kappa_0 \sigma_0 / (1 - \beta)$ that MTP-3's predetermination forces into any closure of (2); the static product $\kappa_0^2 \sigma_0^2$ understates the floor by the factor $(1 - \beta)^{-2} > 1$ and is retired here. Exact CE alone does not restrict h_i^+: with the forward-gap channel open the floor is the sign-free bound (32), never zero. ■

Corollary 4.2 (Forward Guidance Failure)**.** *A forward guidance commitment* $\{R_{t+k}\}_{k=0}^{\infty}$ *does not constitute an infinite-dimensional instrument; the impossibility of Theorem 4.1 binds at every announced horizon.*

Proof. Each announced rate $R_{t+k} \in \mathbb{R}$ is conditioned on $\Omega_{CB,t}$ under capacity constraint $\bar{H}_{CB} < \infty$ (Assumption 2.2); the entire guidance path is a finite-rank projection of the structural state $\theta \in \mathbb{R}^K$, carrying at most $\bar{H}_{CB}$ independent dimensions of information. Since $K \geq 2$, the rank-nullity theorem delivers $\dim(\ker \delta_{t+k}) = K - 1 \geq 1$ at every horizon $k \geq 0$; the floor $A(F, t + k) > 0$ (Lemma 3.9) holds at every finite $t + k$, independently of the guidance horizon. ■

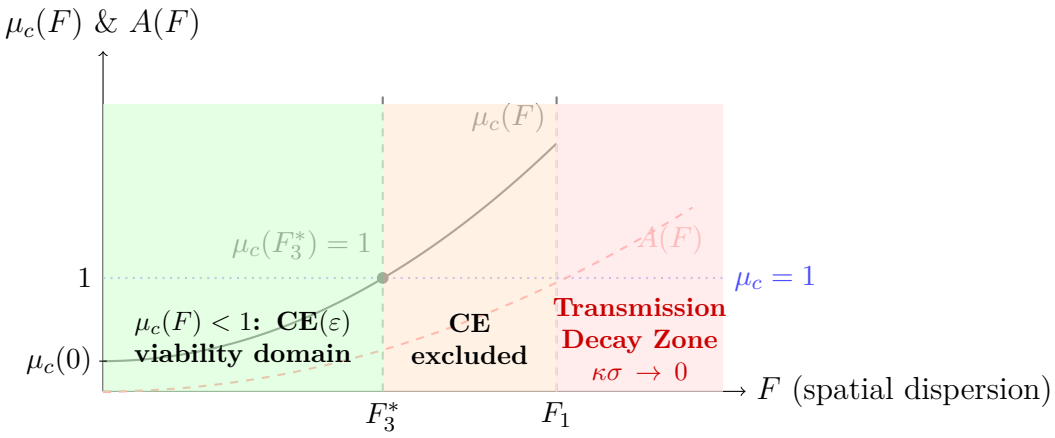


Figure 1: Equilibrium regime partition over spatial dispersion. The solid curve is the expectation-divergence coefficient $\mu_c(F)$; the invariant thresholds F_3^* and F_1 partition the domain. The dashed curve depicts the leading-order approximation $A(F) \approx C_\lambda^2 F^2(\pi^* - \bar{\xi})^2$ to the mechanism-invariant floor under the benchmark affine filtration class; $A(0) = 0$ under perfect spatial homogeneity and the floor grows strictly from the origin at every positive spatial divergence.

5. Threshold Characterization and Structural Limits

5.1 Characterization of the Threshold Boundaries

By Lemma 3.6 and Assumption 2.6, $\kappa(F)\sigma(F)$ is invariant across $[0, F_1)$. Condition (21) imposes $\mu_c(F_1^-) > 1$. These conditions jointly mandate the strict ordering $0 < F_3^* < F_1 < \infty$.

The coordination threshold is the first crossing $F_3^* = \inf\{F \in [0, F_1) : \mu_c(F) \geq 1\}$, so $\mu_c(F_3^*) = 1$ by continuity whenever the infimum is interior. Substituting $\mu_c(F) = a_\pi^2 \bar{X}(F)^2 = (1 - \bar{\lambda}(F))^2 B$, $B \equiv a_\pi^2$, pins the mean non-credibility weight at the boundary:

$$1 - \bar{\lambda}(F_3^*) = B^{-1/2} = \frac{1 - \beta}{\kappa_0 \sigma_0}. \tag{41}$$

This weight is determined by the amplification constant B alone, independently of the specific path of $\bar{\lambda}(F)$.

Proposition 5.1 (Benchmark Coordination Threshold)**.** *Categorical CE(ε) failure holds universally from primitives (P1)–(P3) (Theorem 4.1's Universal Infeasibility). Under the affine credibility-attenuation benchmark, $\bar{\lambda}^{-}(F) = \bar{\lambda}_0 - C'_\lambda F$ with $C'_\lambda > 0$ (satisfying Assumption 3.11), the coordination threshold satisfies:*

$$F_3^* = \frac{\dfrac{1-\beta}{\kappa_0\sigma_0} - (1-\bar{\lambda}_0)}{C'_\lambda}. \tag{42}$$

Proof. The affine path gives $\bar{X}(F) = 1 - \bar{\lambda}_0 + C'_\lambda F$. Substituting into the root condition (23), $1 - \bar{\lambda}_0 + C'_\lambda F_3^* = (1-\beta)/(\kappa_0\sigma_0)$; isolating F_3^* delivers (41) directly, with no quadratic step required since $\mu_c(F)$'s square root is linear in $\bar{X}(F)$. Existence and uniqueness on $(0, F_1)$ follow from Proposition 3.4 under (21), which for the affine path reads $1 - \bar{\lambda}_0 < (1-\beta)/(\kappa_0\sigma_0) < 1 - \bar{\lambda}_0 + C'_\lambda F_1$. ■

The scalars C'_λ and $C^2_\lambda \equiv \lim_{F\to 0^+} \mathrm{Var}_i(\lambda_i)/F^2$ govern the first-order attenuation rate and second-order spatial dispersion of the filtration geometry. The mean decay rate C'_λ determines F_3^* through (41). The variance dispersion rate C^2_λ is a benchmark refinement, not a universal requirement. Lemma 3.7 secures $A(F,t) > 0$ at every finite $F > 0$ through the strict monotonicity $g(F,t) > 0$ alone, independently of the rate of growth. Within the benchmark affine filtration class, where $\bar{\lambda}(F) = \bar{\lambda}_0 - C'_\lambda F$ exactly, the second-order expansion gives $\mathrm{Var}_i(\lambda_i) \sim (C'_\lambda)^2 F^2$ as $F \to 0^+$, so $C^2_\lambda = (C'_\lambda)^2 > 0$ follows directly from $C'_\lambda > 0$ (Proposition 3.4, Uniqueness); the asymptotic floor satisfies $A(F,t) \sim C^2_\lambda F^2(\pi^* - \bar{\xi}(t))^2$ near the origin, giving the quadratic rate of floor growth without conditioning the universal impossibility. The two scalars are structurally distinct and enter different parts of the impossibility architecture.

5.2 Structural Threshold Characterization

Proposition 5.2 (Structural Threshold Determination)**.** (i) Primitive Determination. *The coordinates* $\big(F, \gamma_I, a_\pi\big)$*, equivalently* $(F, \gamma_I, \kappa_0\sigma_0, \beta)$*, are determined by the structural transmission primitives, independent of the expectation-formation mechanism; the superseded combination* $\beta + \kappa_0\sigma_0$ *plays no role in any surviving result.*

(ii) Aggregate Transmission Ceiling. *The capacity-constraint release boundary holds purely analytically:*

$$F_1 \equiv \sup\left\{ F \geq 0 \;\middle|\; \frac{D^+V(h_a)\, h_a}{Kc} > 1 + \gamma_I^2 F^2 \right\}.$$

(iii) Filtration Geometry. *The dispersion constant* $C_\lambda^2 \equiv \lim_{F\to 0^+} \mathrm{Var}_i(\lambda_i)/F^2$ *establishes structurally from the cross-section. The derivation of* F_3^* *requires the right-directional attenuation rate* $C'_\lambda \equiv \limsup_{F\to 0^+}[\bar{\lambda}(0) - \bar{\lambda}(F)]/F > 0$*.* F_3^* *is the exact coordination threshold on the affine tangent of* $\bar{\lambda}(F)$ *at the origin: exact under affine decay and a first-order approximation to the root under nonlinear decay. The aggregate transmission ceiling* F_1 *bounds all institutional comparative statics to the closed-form limits established in item (ii). The floor* $A(F,t) > 0$ *follows from two structurally independent inputs:* $g(F,t) > 0$ *(Lemma 3.7) and* $(\pi^* - \bar{\xi}(t))^2 > 0$ *(scalar non-invertibility, D1–D3); the product is unconditional across every admissible mechanism.*

6. Regime Existence and Threshold Invariance

6.1 Existence of Equilibrium Regimes

Regime I ($F < F_3^*$): CE(ε) fails for all $\varepsilon < A(F,t)$; the finite-horizon dispersion floor $A(F,t) > 0$ binds at every finite t independently of the expectation-divergence coefficient $\mu_c(F)$, which satisfies $\mu_c(F) < 1$ on this interval and does not generate asymptotic expansion.

Regime II ($F < F_3^*$, multi-instrument): Multidimensional fiscal transfers operating outside $\mathfrak{M}(\mathfrak{D})$ absorb the rank-deficiency, neutralizing the unspanned structural variance; equilibrium stabilization requires the fiscal authority to observe θ_i without bandwidth constraint, violating P1.

Regime III ($F_3^* < F < F_1$): Expectation coordination fails permanently via asymptotic divergence, $\mu_c(F) > 1$; PE achieves the bounded floor $\kappa(F)\sigma(F) > 0$ and US admits a positive stabilization floor below F_1. Above F_1, aggregate extraction collapses continuously toward zero, severing policy transmission.

Corollary 6.1 (Regime II Existence)**.** *An equilibrium of type Regime II exists if and only if fiscal capacity offsets the universal floor $A(F,t) > 0$ (Lemma 3.7). The thresholds F_1 and F_3^* are invariant to the number of policy instruments.*

Assumption 6.1 (Loss Divergence — Regime III)**.** For $F > F_3^*$, the effective discounted eigenvalue satisfies $\beta\mu_c(F) \geq 1$.

Remark 6.1 (Relation to Assumption 3.11). Assumption 6.1 requires $\mu_c(F) \geq 1/\beta$, strictly stronger than $\mu_c(F) > 1$ since $\beta \in (0,1)$. Condition (21) delivers $\mu_c(F_1^-) > 1$; Assumption 6.1 is an additional parametric restriction, satisfied whenever $\bar{X}(F_1)^2 \geq (1-\beta)^2/(\beta\kappa_0^2\sigma_0^2)$, since $\mu_c(F) \geq 1/\beta$ reads $a_\pi^2\bar{X}^2 \geq 1/\beta$ exactly. It is a sufficient

condition for the regime comparison in Corollary 6.2, not a consequence of earlier assumptions.

Assumption 6.2 (Regime II Transfer Technology — Benchmark)**.** Regime II supplements the scalar rate with a transfer field removing a fraction $\chi \in [0,1]$ of every region's natural-rate gap q_i from the IS block. Transfers are financed distortionarily and the deadweight cost enters the objective as $\Theta\chi^2$ with $\Theta \in (0,\infty)$. Define the *accommodation ratio*

$$\omega \;\equiv\; \frac{\left(a_\pi^2 + \alpha\sigma_0^2\right)F^2}{\Theta} \;\geq\; 0, \tag{43}$$

the dispersion loss Regime II can remove per unit of fiscal distortion required to remove it. The coefficients are the corrected dynamic loadings: a_π on inflation, σ_0 on the output gap. The informational coefficient γ_I does not appear, and no comparison of the form $\sigma_0 > \gamma_I$ is used.

Corollary 6.2 (Regime Losses and the Non-Existence of an Interior Accommodation Threshold)**.** *F_3^* is strictly necessary for Regime I: for every $F > F_3^*$, Regime I is precluded. Under Assumption 6.1 a finite T^* exists such that the discounted Regime III loss strictly exceeds the transition loss for every $T \geq T^*$, with T^* strictly decreasing in $\mu_c(F) - 1$.*

For $F \leq F_3^*$, under Assumption 3.3 the Regime I regional dispersions are exactly

$$\mathrm{Var}_i(\pi_i - \pi^*) = a_\pi^2\,\mathrm{Var}_i(r_i^*) + \sigma_\varepsilon^2, \qquad \mathrm{Var}_i(y_i - y_i^*) = \sigma_0^2\,\mathrm{Var}_i(r_i^*) + \sigma_\varepsilon^2, \tag{44}$$

so, with $\mathrm{Var}_i(r_i^*) = F^2$ and the common centering term $C_t \equiv (\bar{\pi}_t - \pi^*)^2 + \alpha(\bar{y}_t^{\mathrm{gap}})^2$ suppressed — Assumption 6.2's transfer field removes a fraction of the dispersion *q_i and leaves every aggregate untouched, so C_t enters both regimes identically and*

cancels exactly in (46) —

$$L_{\mathrm{I}} = \left(a_\pi^2 + \alpha\sigma_0^2\right)F^2 + (1+\alpha)\sigma_\varepsilon^2, \tag{45}$$

and, minimizing over χ *under Assumption 6.2,*

$$\chi^* = \frac{\omega}{1+\omega} \in [0,1), \qquad L_{\mathrm{II}} = \Theta\,\frac{\omega}{1+\omega} + (1+\alpha)\sigma_\varepsilon^2. \tag{46}$$

Hence, exactly,

$$L_{\mathrm{I}} - L_{\mathrm{II}} = \Theta\,\frac{\omega^2}{1+\omega} \;\geq\; 0, \tag{47}$$

strictly positive for every $\omega > 0$*, that is for every* $F > 0$ *and every finite* Θ*.* No interior accommodation threshold ω^* exists*: the sign of* $L_{\mathrm{I}} - L_{\mathrm{II}}$ *is determined, not parameterized, and the superseded threshold formula is deleted. The exclusion of Regime II is a feasibility statement, not a welfare one: Regime II requires the fiscal authority to observe* θ_i *region by region without bandwidth constraint, violating (P1), and it therefore leaves* $\mathfrak{M}(\mathfrak{D})$ *by Condition (C) of Definition 2.2. The thresholds* F_1 *and* F_3^* *remain invariant to the number of instruments (Corollary 6.1).*

Proof. Necessity of F_3^* for Regime I and the T^* statement follow from $\mu_c(F) > 1$ above the first crossing (Proposition 3.4(iii)) with Assumption 6.1. For (43): under Assumption 3.3 ($h_i^+ = 0,\ f_i = 0$), Lemma 4.1 gives $\pi_{it} - \bar{\pi}_t = a_\pi q_i$, and (3) with R_t common across i and $\mathbb{E}_{i,t}[\pi_{i,t+1}]$ undispersed gives $y_{it}^{\mathrm{gap}} - \bar{y}_t^{\mathrm{gap}} = \sigma_0 q_i$; adding the idiosyncratic residual of variance σ_ε^2 (D5), independent of q_i by Assumption 2.1, gives the two displayed dispersions. The loss (4) is mean-square, so it decomposes as $\int_0^1(\pi_{it} - \pi^*)^2 di = \mathrm{Var}_i(\pi_{it}) + (\bar{\pi}_t - \pi^*)^2$ and $\int_0^1(y_{it}^{\mathrm{gap}})^2 di = \mathrm{Var}_i(y_{it}^{\mathrm{gap}}) + (\bar{y}_t^{\mathrm{gap}})^2$: the centering term C_t is not zero and is not claimed zero from Assumption 3.3, which restricts forecast dispersion and says nothing about aggregates. It is common to the

two regimes and cancels in the difference, giving (44) up to C_t. Under Assumption 6.2 the residual gap is $(1-\chi)q_i$, so

$$L_{\mathrm{II}}(\chi) = \left(a_\pi^2 + \alpha\sigma_0^2\right)(1-\chi)^2F^2 + (1+\alpha)\sigma_\varepsilon^2 + \Theta\chi^2,$$

strictly convex in χ, with first-order condition $-(1-\chi)\Theta\omega + \Theta\chi = 0$ by (42), hence $\chi^* = \omega/(1+\omega)$ and $L_{\mathrm{II}} = \Theta\omega/(1+\omega)^2 + \Theta\omega^2/(1+\omega)^2 + (1+\alpha)\sigma_\varepsilon^2 = \Theta\omega/(1+\omega) + (1+\alpha)\sigma_\varepsilon^2$. Subtracting from (44), and using $\left(a_\pi^2 + \alpha\sigma_0^2\right)F^2 = \Theta\omega$, gives $L_{\mathrm{I}} - L_{\mathrm{II}} = \Theta\omega - \Theta\omega/(1+\omega) = \Theta\omega^2/(1+\omega)$, which is (46). Every branch is covered: at $\omega = 0$, that is at $F = 0$ or $\Theta = \infty$, the two regimes tie; at every $\omega > 0$ Regime II strictly dominates. No branch requires $a_\pi^2 \operatorname{Var}_i(r_i^*) > \sigma_\varepsilon^2$, and no denominator can vanish, since $1 + \omega \geq 1$. ■

Remark 6.2 (What the Retired ω^* Concealed). The superseded threshold compared Regime I with an undisplayed Regime II loss and used the static coefficient $\gamma_A = \sigma_0/(1+\kappa_0\sigma_0)$ inside a dynamic cross-section, where Lemma 3.1 requires a_π on inflation and σ_0 on the output gap; its nonsingularity condition $a_\pi^2 \operatorname{Var}_i(r_i^*) > \sigma_\varepsilon^2$ protected a denominator that the corrected comparison does not contain. Once both losses are displayed and both coefficients corrected, the comparison has no interior root. The economic content of Regime II is unchanged and sharpened: it is excluded by the capacity primitive, never by welfare.

6.2 Structural Bounds on Expectation Formation

Proposition 6.1 (Comparative Statics and the Supremum Bound)**.** *The equilibrium thresholds F_1 and F_3^* are strictly invariant to any primitive that preserves the parameter vector $(\kappa_0, \sigma_0, h_a, K, c, C_\lambda^2, C_\lambda')$. An exogenous increase in h_a implies $dF_1 >$*

0*; a reduction in* C'_λ *implies* $dF_3^* > 0$ *via* (41)*; a reduction in* c *implies* $dF_1 > 0$.

The aggregate transmission ceiling bounding expectation stability satisfies

$$\Phi = \sup\{F : \mu_c(F, h_a^{\max}, C'^{\min}_\lambda) < 1\} < \infty \tag{48}$$

for every mechanism in $\mathfrak{M}(\mathfrak{D})$*. This bound is finite at the best feasible institutional design: expectation coordination fractures at a strictly finite fragmentation distance determined by the jurisdiction's structural primitives, not by its institutional choices.*

6.3 Threshold Properties and Parameter Restrictions

Shifts in structural primitives displace these threshold coordinates; the underlying rank constraints are invariant.

Regime I ($F < F_3^*$): F_3^* is the first crossing above which benchmark forecast dispersion diverges; below it dispersion is bounded, which is not exact coordination, since Theorem 4.1 denies CE(0) at every $F > 0$. An exogenous increase in h_a shifts F_1 upward, expanding the parameter space supporting Regime III; F_3^* remains invariant.

Regime II ($F < F_3^*$): By Corollary 6.2, $L_{\mathrm{I}} - L_{\mathrm{II}} = \Theta\omega^2/(1+\omega) \geq 0$ on the admissible domain, with equality only at $\omega = 0$: the multi-instrument parameterization weakly dominates the single-instrument benchmark on welfare at every $\omega > 0$, and no interior accommodation threshold ω^* exists. Regime II is excluded by the capacity primitive (P1), never by the welfare ordering.

Regime III ($F > F_3^*$): Reductions in $\sigma_F^2(0)$ are dominated by $\mu_c(F) > 1$; asymptotics are invariant to initial conditions. Restoring $\mu_c(F) < 1$ requires either a reduction in C'_λ that elevates F_3^* via (41), or a direct reduction in F.

Corollary 6.3 (Endogenous Regime Fissures)**.** *Divergent equilibrium outcomes across disjoint jurisdictions arise as optimal boundary solutions when jurisdictions*

locate on opposite sides of F_3^*. Because the feasibility set $\mathcal{F}(F)$ contracts discontinuously at F_3^*, divergent macroeconomic regimes represent geometric adaptations to differing limits on expectation feedback, not variations in institutional preference.

The instability condition for $F > F_3^*$ governs universally any expectation-formation mapping satisfying structural non-concentration (Assumption 2.3) under (P1)–(P3).

7. Conclusion

The analytical failure of the classical macroeconomic trilemma structurally dictates the operational limits of the central bank. Because a scalar interest rate cannot span multidimensional regional structural divergence, communication frameworks and institutional credibility mechanisms cannot eliminate the resulting persistent spatial fragmentation.

Cross-sectional dispersion in expectations and localized inflation outcomes is not a consequence of institutional failure. It follows from the dimensional structure of the problem: multi-dimensional structural heterogeneity maps through a scalar instrument, and the unspanned residual cannot be eliminated by design.

Every admissible mechanism operates within limits that institutional design cannot remove: one instrument confronts $K \geq 2$ dimensions of structural variation. The feasibility set $\mathcal{F}(F)$ contracts monotonically in spatial divergence; no policy rule, no institutional architecture, and no expansion of instrument count eliminates this contraction. Expectation coordination, policy transmission, and uniform regional stabilization cannot be achieved simultaneously because a scalar instrument cannot reach a multidimensional target—this failure is dimensional, not behavioral.

The impossibility is not a policy failure.

It is the shape of the world.

The central bank did not build the cage.

It is operating inside one.

References

Alexandrov, A., and M. K. Brunnermeier (2026). Optimal (un)conventional monetary policy. Working paper, Princeton University.

Angeletos, G.-M., and C. Lian (2018). Forward guidance without common knowledge. *American Economic Review* 108(9): 2477–2512.

Arrow, K. J. (1951). *Social Choice and Individual Values*. New Haven: Yale University Press.

Auclert, A. (2019). Monetary policy and the redistribution channel. *American Economic Review* 109(6): 2333–2367.

Barro, R. J., and D. B. Gordon (1983). A positive theory of monetary policy in a natural rate model. *Journal of Political Economy* 91(4): 589–610.

Calvo, G. A. (1983). Staggered prices in a utility-maximizing framework. *Journal of Monetary Economics* 12(3): 383–398.

Clarida, R., J. Galí, and M. Gertler (1999). The science of monetary policy. *Journal of Economic Literature* 37(4): 1661–1707.

Cúrdia, V., and M. Woodford (2010). Credit frictions and optimal monetary policy. *Journal of Money, Credit and Banking* 42(S1): 3–35.

Cover, T. M., and J. A. Thomas (1991). *Elements of Information Theory*. New York: Wiley.

Friedman, M. (1968). The role of monetary policy. *American Economic Review* 58(1): 1–17.

Hottman, C., and R. Reyes-Heroles (2024). Globalization, inflation dynamics, and the slope of the Phillips curve. Working paper, Federal Reserve Board.

Judd, K. L. (1985). The law of large numbers with a continuum of IID random variables. *Journal of Economic Theory* 35(1): 19–25.

Kaplan, G., B. Moll, and G. L. Violante (2018). Monetary policy according to HANK. *American Economic Review* 108(3): 697–743.

Kenen, P. B. (1969). The theory of optimum currency areas: An eclectic view. In R. Mundell and A. Swoboda, eds., *Monetary Problems of the International Economy*. Chicago: University of Chicago Press.

Lucas, R. E. (1976). Econometric policy evaluation: A critique. *Carnegie-Rochester Conference Series on Public Policy* 1: 19–46.

Mackowiak, B., and M. Wiederholt (2009). Optimal sticky prices under rational inattention. *American Economic Review* 99(3): 769–803.

McKinnon, R. I. (1963). Optimum currency areas. *American Economic Review* 53(4): 717–725.

Morris, S., and H. S. Shin (2002). Social value of public information. *American Economic Review* 92(5): 1521–1534.

Mundell, R. A. (1961). A theory of optimum currency areas. *American Economic Review* 51(4): 657–665.

Sims, C. A. (2003). Implications of rational inattention. *Journal of Monetary Economics* 50(3): 665–690.

Stokey, N. L., R. E. Lucas, and E. C. Prescott (1989). *Recursive Methods in Economic Dynamics*. Harvard University Press.

Woodford, M. (2003). *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton University Press.