

Toward Formalizing Nash's Ideal Money: A Robust Framework for Preserving Price Stability, and Incentive Compatibility Through Index Anchoring and Resilience Objectives in Monetary Policy

[J. G. B. Northey](#) and [M. P. Ivy](#) *

Posted Date: 3 February 2026

doi: 10.20944/preprints202601.1981.v2

Keywords: ideal money (Nash); index anchoring; robust control; price-level targeting; moral hazard; term-premium wedge; safe-set invariance; macro-financial resilience



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

Toward Formalizing Nash’s Ideal Money: A Robust Framework for Preserving Price Stability, and Incentive Compatibility Through Index Anchoring and Resilience Objectives in Monetary Policy

J. G. B. Northey and M. P. Ivy *

Frontier Agri-Science Inc., Agricultural Policy Branch
* Correspondence: michael.ivy1@gmail.com

Abstract

We develop a formal theory of “ideal money” that extends Nash’s proposal for a nonpolitical value standard and his critique of discretionary monetary authorities as implicit “pardoners” of overindebtedness. The core contribution is a precise, equilibrium-compatible notion of minimal distortion to productive incentives: a sound regime should preserve the unit of account for long-horizon contracting while limiting the incentive wedges created by nominal risk, expected rescues, and targeted price suppression. We cast monetary design as a robust feedback-control problem around an objective index anchor, defining a tracking error e_t and establishing bounded-gain stability under bounded (potentially adversarial) disturbances. We then broaden “soundness” from scalar anchor tracking to system-level resilience by requiring safe-set invariance of a macro-financial state that includes inflation uncertainty, leverage/maturity mismatch, and credit stress. To formalize modern interventions, we introduce explicit policy wedges: a relative-price wedge τ_t and a pardon/transfer wedge b_t , and we show how these wedges propagate through collateral and balance-sheet dynamics into credit spreads and sectoral tilts in capital allocation and labor deployment. Finally, we adopt a two-stage treatment of interventions: exogenous “distortion accounting” to quantify the induced misallocation for any bounded wedge path, and an endogenous political-economy layer in which constitutional constraints reduce the equilibrium frequency and magnitude of such wedges. Together, these elements provide a mathematically explicit criterion for sound money: a regime that robustly preserves the value standard, remains resilient under shocks, and minimizes measurable incentive distortions induced by discretion.

Keywords: ideal money (Nash); index anchoring; robust control; price-level targeting; moral hazard; term-premium wedge; safe-set invariance; macro-financial resilience

1. Background And Economic Problem

The practical value of money is not exhausted by its role as a medium of exchange. Money is also a unit of account for long-horizon commitments: wages, leases, mortgages, corporate debt, and the intertemporal planning of households and firms. When the value of the unit is unreliable, agents rationally shorten horizons, demand risk premia, and allocate resources to hedging and political gaming rather than to productive activity.

Price stability as preservation of the unit of account. In contemporary terms, the core monetary objective may be stated as preserving the real value of the monetary unit over planning-relevant horizons: the unit of account should remain sufficiently stable that nominal contracts map predictably into real outcomes. In this paper we operationalize “price stability” not as a slogan, but as a measurable control objective: deviations of the domestic price level P_t from an objective benchmark index I_t (with an optional design drift π^*) should be small, mean-reverting, and robust to shocks. This connects the practical mandate of inflation-adjusted value preservation directly to the tracking-error variable e_t and to the resilience conditions introduced below.

1.1. From Phillips To Lucas: The Limits Of Exploitable Stabilization Tradeoffs

Mid-century macroeconomic policy debates were strongly shaped by the empirical relationship documented by Phillips [1] and its subsequent interpretation as a usable menu of inflation–unemployment outcomes. In influential postwar discussions, the Phillips relationship was read as suggesting a policy-relevant tradeoff between inflation and unemployment, consistent with a broader Keynesian stabilization ambition of managing aggregate demand to improve real outcomes [2,3]. However, the central practical difficulty is that the apparent tradeoff is neither structurally invariant nor reliably exploitable over time: once private agents revise expectations and contract terms in response to systematic policy behavior, the mapping from policy actions to real outcomes changes.

This critique was articulated in two steps. First, Friedman [4] argued that attempts to maintain unemployment below its sustainable (“natural”) level by tolerating persistently higher inflation would fail as expectations adjust, leaving only higher inflation and degraded price-level predictability. Second, Lucas [5] showed that econometric relations estimated under one policy regime cannot be used to evaluate alternative regimes because agents’ decision rules (and thus reduced-form parameters) change when the policy rule changes. In combination, these insights shifted the design problem away from discretionary “fine-tuning” and toward credible, systematic monetary rules that stabilize expectations and preserve the unit of account.

1

1.2. From Lucas To Nash: From Rules-Of-Thumb To A Verifiable Monetary Constitution

The Lucas critique does not by itself specify *which* rule should be adopted; it implies that the rule must be explicit, credible, and evaluated as a regime. This motivation is closely related to the classic time-inconsistency problem in monetary policy: absent credible commitment, discretionary stabilization incentives can systematically erode the credibility of the nominal anchor [7,8]. One influential practical response is the class of simple feedback rules in which the policy instrument reacts systematically to observable deviations from targets, exemplified by Taylor-type interest-rate rules [9]. Such rules aim to reduce discretion by making policy systematic and therefore learnable by private agents.

John Nash’s *Ideal Money* pushes the rules agenda in a distinct direction: it argues that the *unit of account itself* should be constitutionally anchored to a nonpolitical value standard—an externally referenced index built from internationally traded commodity prices—and that the construction, smoothing, and revision of this standard must be governed by rule-bound protocols insulated from opportunistic redefinition [10]. Nash’s motivation is explicitly contractual and incentive-oriented: when authorities retain broad latitude to debase, re-anchor, or provide open-ended relief, they operate as “grand pardoners” of overindebtedness in a way that private agents anticipate *ex ante*, thereby distorting leverage, maturity choice, and risk-taking [10]. The design implication is a *monetary constitution*: a mechanically computable standard, a published refinement methodology, and explicit bounds on the discretionary interventions that create predictable wedges in relative prices and bailout expectations [10].

1.3. This Paper: Ideal Money As Robust Index Anchoring With Bounded Discretion

This paper formalizes Nash’s program as an implementable, verifiable monetary design with three linked aims: (i) preserve the unit of account for long-horizon contracting; (ii) limit incentive distortions generated by expected rescues and targeted relative-price interventions; and (iii) maintain macro-financial resilience under stress.

Concretely, we contribute:

¹ This paper treats “employment stabilization mandates” as a historically contingent design choice rather than a primitive. The Federal Reserve’s founding statute (1913) did not specify a maximum-employment objective; explicit macroeconomic goals including “maximum employment” and “stable prices” were made statutory later (1977 amendments) [6].

1. **A reproducible nonpolitical anchor.** We define a transparent international commodity anchor and its published (smoothed/refined) form, together with an explicit constitutional revision protocol [10].
2. **A regime-level soundness criterion.** We define robust soundness as a bounded-gain property for the price-level gap e_t and derive exact, checkable bounds under a simple error-correction recursion.
3. **A contracting-relevant nominal-risk statistic.** We measure nominal instability at horizon H via a cumulative relative-inflation variance $\sigma_{t,H}^2$, aligning the criterion with the maturities that govern real investment and credit-market structure [10].
4. **Explicit representation of discretion as wedges.** We represent modern interventions as measurable wedges: a relative-price wedge τ_t and a pardon/transfer-capacity wedge b_t , and we show how bounding these objects reduces predictable moral hazard and misallocation.
5. **Resilience as an engineering constraint.** We extend soundness from scalar anchor tracking to system-level resilience by imposing safe-set invariance and providing verifiable sufficient conditions and trajectory bounds.

Our approach is closely related to three established lines of monetary research, but differs in what it makes *verifiable* as an institutional design. First, the rules-versus-discretion tradition that follows the Lucas critique motivates policy rules that are credible as regimes (rather than one-off optimizations), including simple, implementable feedback rules such as Taylor-type instrument rules and their modern New Keynesian implementations [5,9,11]. Second, robust-control approaches to monetary policy under model misspecification motivate treating “soundness” as a bounded-gain property and designing policies that remain stable under bounded disturbances [12]. Third, work on crisis policy and moral hazard emphasizes that anticipated support can raise ex ante risk-taking and distort allocation even when short-run stabilization succeeds [13]. Separately, price-level (path) targeting highlights the value of history dependence for long-run anchoring of the price level [14–16]. The distinctive contribution here is to integrate these themes into a *monetary constitution* in Nash’s sense: a nonpolitical, externally referenced value standard with explicit construction and revision governance [10], combined with (i) provable bounded-gain tracking of the level gap, (ii) contracting-relevant horizon risk, (iii) explicit representation and constitutional caps on intervention wedges, and (iv) a system-level resilience requirement stated as safe-set invariance.²

1.4. Terminology: Rules, Discretion, And Disturbances

We distinguish three concepts that are often conflated in monetary commentary:

1. **Rule-based policy actions:** instrument adjustments that are pre-specified as a function of public observables (e.g., a feedback rule that reacts to e_t within a known tolerance band).
2. **Discretionary actions:** interventions not implied by the published rule or not bounded by its pre-committed parameters (e.g., targeted asset purchases that suppress specific yields, ad hoc rescues, emergency programs whose scale is not constrained ex ante). Discretion is modeled either as deviations from the rule or as additional policy wedges that directly alter relative prices.
3. **Disturbances:** exogenous or imperfectly controlled influences that enter the closed-loop dynamics.

In the framework below, disturbances enter as bounded sequences (e.g., w_t in the tracking recursion and ξ_t in the macro-financial state dynamics). Discretion enters as explicit wedges (τ_t, b_t) that are constitutionally bounded and explicitly penalized in the distortion criterion.

2. Nash’s Core Vision as Formal Design Requirements

Nash’s paper contains three elements that translate directly into formal design constraints for monetary policy:

² Related normative framing appears in policy debate: Shelton characterizes money as “a moral contract” and stresses that it should function as a uniform measure that “has to work the same for everyone” [17].

1. **Long-horizon contracting reliability.** Money quality matters most for contracts with a long time axis [10]. This points to stabilizing not merely one-step inflation, but the uncertainty in the future price level over contract maturities.
2. **Nonpolitical value standard.** Nash’s ICPI proposal is explicitly based on international commodity prices and aims to remove the political role of “grand pardoners” [10]. This requires an explicit index definition and governance rules for revisions.
3. **Refined indices and smoothing.** Nash argues that an index can be designed to vary smoothly in the short run yet follow industrial-input costs over longer horizons via moving averages and “pricing modulo” a smoother sub-index [10]. This motivates an explicit filter for the anchor and a transparent estimation policy.

Two additional Nash considerations are also central for policy optimization:

1. **Political corruption risk.** Nash notes that a political basis for changing the index invites corruption; a constitutional protocol should make index tampering as unlikely as tampering with the meter or kilogram [10]. This is a governance constraint, not merely a statistical preference.
2. **“Safe-deposit box” singularity.** Nash observes that if money is “too good” as a store of value, it can become an unpriced global safe-deposit facility; a completely steady and constant inflation rate can avoid this while preserving predictability (contractible inflation) [10]. This motivates allowing a design drift π^* rather than insisting on $\pi^* = 0$.

Table 1 summarizes the Nash-derived design requirements and the specific formal objects used to operationalize each requirement.

Table 1. Nash-derived design requirements and operational counterparts

Requirement	Why it matters (economic intuition)	Operationalization in this paper
Long-horizon contracting reliability	Money quality is revealed in long-dated credit: price-level uncertainty raises premia and shortens maturities.	Contracting-risk statistic $\sigma_{t,H}^2$ (Def. 5); level gap e_t (Def. 3).
Nonpolitical value standard (ICPI logic)	An internationally referenced anchor reduces scope for domestic discretionary redefinition of the standard.	Published anchor construction I_t (Def. 2); raw index $\hat{I}_t^*$ (Eq. (1)).
Refined indices and smoothing	Commodity prices can be informative but volatile; refinement improves implementability by filtering transitory noise.	Baseline smoothing filter (Eq. (3)); refined-index variant (App. D).
Political corruption risk (governance)	If index rules can be altered ad hoc, the standard becomes a political instrument; credibility requires a rule-bound protocol.	Constitutional revision protocol (Section 3.4) embedded in Def. 2.
“Safe-deposit box” singularity / drift choice	Money that is “too good” as a store of value can attract unpriced global hoarding; a contractible drift avoids this.	Design drift π^* in the value standard (Def. 3).

Here we define a transparent, externally referenced anchor; we stabilize the *level* of the price standard relative to that anchor; we measure and penalize *horizon* nominal risk; we represent discretionary relief and targeted price manipulation as explicit wedges; and we impose resilience constraints that reduce crisis-driven demand for discretionary pardons.

3. The Nonpolitical Value Standard: Index Construction, Smoothing, Governance

3.1. Raw ICPI in a Reference Numeraire

Let $p_{i,t}^*$ denote the international price of commodity $i \in \{1, \dots, N\}$ at time t quoted in a reference numeraire (e.g., USD or SDR). Let weights $w_i \geq 0$ satisfy $\sum_i w_i = 1$. Define the raw international ICPI (geometric):

$$\log \hat{I}_t^* := \sum_{i=1}^N w_i \log p_{i,t}^*. \quad (1)$$

Rationale. Quoting component prices in a reference numeraire makes the anchor internationally comparable and minimizes dependence on domestic policy choices. The log-geometric aggregator delivers additivity: changes in $\log \hat{I}_t^*$ are weighted averages of component log price changes, simplifying smoothing and the derivation of relative-inflation control rules. Fixing the basket and weights makes the construction mechanically reproducible, which is essential for a nonpolitical standard and for long-horizon contracting reliability [10].

This form is scale-free and makes the anchor interpretable as a weighted average of log commodity prices.

3.2. Domestic-Currency Anchor and Numeraire Choice (Option A vs Option B)

Let E_t denote the nominal exchange rate (domestic currency per unit of reference numeraire). A domestic-currency representation of the raw index is:

$$\hat{I}_t := E_t \hat{I}_t^*, \quad \log \hat{I}_t = \log E_t + \log \hat{I}_t^*. \quad (2)$$

Remark 1 (FX conversion and the geometric aggregator). *If all component prices are quoted in the same reference numeraire, then converting each component price to domestic currency before aggregation yields the same raw index as aggregating first and then multiplying by E_t , because $\sum_i w_i \log(E_t p_{i,t}^*) = \log E_t + \sum_i w_i \log p_{i,t}^*$. Thus FX volatility does not require an “adjustment” to the geometric mean; it enters only through the regime choice (Option A versus Option B) regarding whether domestic prices are stabilized in the reference numeraire or in domestic currency.*

Because E_t responds to domestic monetary expectations, the anchor in domestic currency embeds market-driven endogeneity. This is not “political,” but it is a design choice that must be explicit.

Definition 1 (Numeraire design choice). *Two internally consistent anchoring regimes are:*

1. **Option A (global-numeraire anchoring):** define $P_t^* := P_t / E_t$ and anchor P_t^* to $\hat{I}_t^*$ (so policy targets domestic prices in the reference numeraire).
2. **Option B (domestic-currency anchoring):** anchor domestic P_t directly to the domestic-currency index $\hat{I}_t$ defined in (2).

Rationale. The anchor is constructed in a reference numeraire to separate internationally traded price movements from domestic measurement conventions. Converting via the market exchange rate E_t yields the domestic-currency anchor faced by residents. Because E_t is endogenous to expectations (and therefore to policy), it is essential to make the numeraire choice explicit: one may target domestic prices in the reference unit (Option A) or target the domestic-currency anchor (Option B). Stating the choice up front prevents implicit conflation of global commodity-price variation with domestic monetary feedback through FX [10].

Option A cleanly separates domestic policy from FX conversion by working in a global unit; Option B anchors domestic prices to a mechanically computed, market-converted commodity index. Both can satisfy Nash's "nonpolitical standard" criterion if index construction and revision remain rule-based and transparent [10].

3.3. Smoothing And "Refined Indices"

Nash emphasizes that the index serving as a value standard should be *refined* to vary smoothly in the short run yet appropriately over longer periods [10]. We operationalize this by defining the *published* anchor I_t as a smoothed version of $\hat{I}_t$.

Baseline smoothing (exponential filter in logs).

Let $\lambda \in (0, 1]$. Define

$$\log I_t = (1 - \lambda) \log I_{t-1} + \lambda \log \hat{I}_t. \quad (3)$$

Rationale. Commodity prices can track long-run industrial costs while remaining volatile at business-cycle frequencies. The published anchor should therefore be a *filtered* object: smoothing reduces high-frequency noise that would otherwise enter the feedback rule as measurement disturbance. This improves implementability and robustness by reducing instrument overreaction and keeping corrections focused on persistent deviations in the value standard. The refined-index variant implements Nash's idea of combining short-run smoothness from a stable basket with long-run tracking from a more informative basket ("pricing modulo"), using moving averages to separate transitory volatility from long-horizon signal [10].

Smaller λ yields more short-run smoothness but slower response; larger λ yields a more responsive anchor.

Refined-index variant ("pricing modulo" a smoother sub-index).

Nash suggests that one can obtain the smoothness of a "smoother" basket while following the long-run behavior of a more "industrial" basket by applying moving averages to the ratio between them [10]. A formal version is given in Appendix D. The main text uses (3) for clarity.

3.4. Constitutional Revision Protocol

A nonpolitical standard is not only a formula; it requires a rule for *revisions*. We propose a constitutional protocol:

1. **Eligibility:** components must be internationally traded, liquid, and transparently priced.
2. **Weights:** fixed for long windows (e.g., K years); changes occur only at scheduled review dates.
3. **Revision triggers:** objective criteria (liquidity failure, market discontinuation, persistent manipulation risk).
4. **Revision method:** rule-based (e.g., minimize historical short-run volatility subject to industrial relevance constraints).
5. **Publication:** algorithm, data sources, and change logs published; no ad hoc changes.

This is directly responsive to Nash's warning that political control over index changes invites corruption [10].

Definition 2 (Published nonpolitical anchor). *Given internationally traded prices $\{p_{i,t}^*\}_{i=1}^N$ and weights $\{w_i\}_{i=1}^N$ with $w_i \geq 0$ and $\sum_i w_i = 1$, the raw international index is $\log \hat{I}_t^* := \sum_{i=1}^N w_i \log p_{i,t}^*$. Let E_t be the market exchange rate (domestic per reference numeraire) and define $\hat{I}_t := E_t \hat{I}_t^*$. The published anchor I_t is the smoothed index defined by the published filter (e.g., (3)) and is revised only under the constitutional revision protocol.*

3.5. Illustrative Numerical Example And Data Sources

Five-period illustration (raw vs. smoothed anchor).

To make the smoothing rule operationally transparent, Table 2 simulates a five-period raw ICPI and its exponentially smoothed counterpart using hypothetical commodity prices. We use three commodities with weights $(w_{Cu}, w_{Al}, w_{Oil}) = (0.50, 0.30, 0.20)$ and compute the raw index via the geometric aggregator in (1). The smoothed series applies the log-filter in (3) with $\lambda = 0.30$ and is initialized at $I_1^* = \hat{I}_1^*$.

Table 2. Illustrative ICPI construction and smoothing (hypothetical data; indices normalized at $t = 1$)

t	$p_{Cu,t}^*$	$p_{Al,t}^*$	$p_{Oil,t}^*$	$\hat{I}_t^*$ (raw)	I_t^* (smoothed, $\lambda = 0.30$)
1	100.00	50.00	70.00	100.00	100.00
2	110.00	52.00	80.00	108.99	102.62
3	90.00	48.00	60.00	90.87	98.94
4	105.00	51.00	75.00	104.51	100.58
5	95.00	49.00	65.00	95.45	99.01

Notes: Indices are computed in log space for the geometric mean and smoothing, then exponentiated and scaled so that both series equal 100 at $t = 1$ for comparability. Author’s calculations with hypothetical data.

Interpretation. The raw index inherits the full high-frequency volatility of the component prices, whereas the smoothed anchor filters transitory movements and responds primarily to persistent shifts. In this toy example the raw series moves from 100 to 108.99 to 90.87 to 104.51 to 95.45, while the smoothed series remains in a much tighter band around the baseline. This reduction in high-frequency noise supports robust policy implementation by reducing noise in the tracking error e_t and tightening the bounded-gain guarantees established in Proposition 1. A real-world analog can be produced by applying the same construction to publicly available series (e.g., IMF copper/aluminum/crude prices over a short window); one typically obtains a qualitatively similar volatility reduction, and computations can be provided upon request.

Potential data sources (practical implementation).

In an empirical implementation, each $p_{i,t}^*$ can be taken from publicly available commodity price series. For monthly-frequency construction, widely used sources include the IMF *Primary Commodity Prices* database and the World Bank *Commodity Markets* (“Pink Sheet”) series. Both series are public and replication-friendly benchmarks for internationally comparable commodity pricing inputs [18,19]. For higher-frequency construction, one can rely on exchange-quoted spot or nearby-futures prices from major venues (e.g., LME/COMEX for metals and NYMEX/ICE for energy) with a constitutionally specified roll convention and unit standardization. To mitigate manipulation risk (a central Nash concern), the protocol should prioritize multi-source verification, pre-specified outlier rules, and rule-based imputation for missing data (e.g., deterministic carry-forward or median-of-sources rules), with all such rules published ex ante as part of the index constitution [10].

4. Core Objects: Index-Anchored Standard, Tracking Error, and Robust Soundness

Let P_t denote the aggregate domestic price level (in domestic currency units). Let I_t denote the *published* (smoothed) index anchor defined above.

Definition 3 (Index-anchored value standard). *Fix a drift parameter $\pi^* \in \mathbb{R}$ and an initial parity $\kappa > 0$. The target value standard is*

$$\log\left(\frac{P_t}{I_t}\right) = \log \kappa + \pi^* t.$$

(4)

Define the tracking error

$$e_t := \log\left(\frac{P_t}{I_t}\right) - (\log \kappa + \pi^* t). \quad (5)$$

Rationale. The ratio P_t/I_t measures the domestic price level in units of the nonpolitical anchor. The constant κ fixes the initial parity, while π^* allows an explicit, contractible long-run drift (including $\pi^* = 0$). The tracking error e_t is the sufficient statistic for policy design: it is the *level* gap that must be mean-reverting if the unit of account is to remain reliable. Targeting a level gap, rather than only next-period inflation, forces re-anchoring after shocks and is therefore aligned with Nash's emphasis on long-horizon contracting quality as the key test of money [10].

Remark 2 (Nash's "too good money" point and the role of π^*). *Allowing $\pi^* > 0$ admits a constant, contractible drift that can mitigate Nash's "safe-deposit box" singularity: money can remain predictable for contracting while not offering a free global storage technology [10].*

Definition 4 (Robust soundness (bounded-gain form)). *A monetary regime is robustly sound if there exist constants $C < \infty$ and $0 < \rho < 1$ such that for all bounded disturbance sequences $\{w_t\}$,*

$$|e_t| \leq C\rho^t |e_0| + C \sup_{0 \leq s < t} |w_s| \quad \forall t \geq 0. \quad (6)$$

Rationale. The bounded-gain inequality formalizes "reliability" of the unit of account as a robust control property: initial misalignment decays geometrically, and any bounded sequence of shocks produces only bounded deviations. The constants ρ and C summarize the regime's speed of correction and sensitivity to disturbances, respectively. This is the appropriate criterion when the environment includes measurement error, money-demand shifts, and bounded implementation slippage: one does not require perfect stabilization, only that errors cannot drift or explode under plausible stress.

Interpreting the robust-soundness bound. Inequality (6) is a quantitative statement that the index-tracking error is robustly controlled in the presence of shocks. The constants $0 < \rho < 1$ and $C \geq 1$ are determined by the closed-loop dynamics: ρ governs how rapidly deviations are corrected, while C scales sensitivity to disturbances. The inequality decomposes the error into two parts: a decaying "memory" of initial misalignment and a component proportional to the worst disturbance magnitude experienced.

4.1. A Key Identity: Relative Inflation and the Level Gap

Define inflation $\pi_{t+1} := \log P_{t+1} - \log P_t$ and anchor inflation $\pi_{t+1}^I := \log I_{t+1} - \log I_t$. From (5),

$$e_{t+1} - e_t = (\pi_{t+1} - \pi_{t+1}^I) - \pi^*. \quad (7)$$

Thus stabilizing the *level* gap e_t is equivalent to controlling *relative inflation* $\pi_{t+1} - \pi_{t+1}^I$ around π^* .

5. Policy Implementation: A Robust Error-Correction Rule With an Implementable Instrument

5.1. Why an Implementable Instrument Matters

A monetary constitution must be stated in terms of policy instruments that can actually be implemented and audited in real time. In contemporary operating frameworks, broad monetary aggregates are largely endogenous to credit creation and portfolio shifts, whereas the policy rate (and, when relevant, a reserve/settlement-rate operating target) is directly set by the central bank and is therefore the natural instrument for a rule-based regime.

To retain the paper's central robust-control logic while aligning with practice, we proceed in two layers:

1. **Main text:** a rate-based error-correction condition that closes the loop on the value-standard gap e_t and yields exact bounded-gain guarantees.
2. **Appendix:** an equivalent base-control representation showing how a policy-controlled nominal liability can implement the same closed-loop error recursion under standard money-demand relations.

Rationale. The key object for design is the closed-loop recursion for the level gap e_t ; multiple operational implementations can generate that recursion. Stating the constitution in terms of an implementable instrument makes the regime verifiable and prevents “rule drift” from being reintroduced through unobservable or ill-defined operating targets.

5.2. Reduced-Form Error-Correction Condition

Assume the policy regime can implement the following *error-correction* relation:

$$(\pi_{t+1} - \pi_{t+1}^I) - \pi^* = -\phi e_t + w_t, \quad \phi > 0, \quad (8)$$

where w_t aggregates bounded shocks (money demand shifts, measurement noise in P_t or I_t , transitory supply shocks, imperfect neutral-rate estimates, and bounded implementation slippage).

Combining (7) and (8) yields the closed-loop recursion:

$$e_{t+1} = (1 - \phi)e_t + w_t. \quad (9)$$

Rationale. This recursion isolates the closed-loop behavior of the value-standard gap. The gain ϕ governs mean reversion: $\phi \in (0, 2)$ implies $|1 - \phi| < 1$ and therefore geometric correction of deviations, while larger ϕ increases responsiveness but amplifies measurement and transmission noise in w_t . Writing the system in this AR(1) form makes robustness transparent and yields exact bounds used to set policy tolerances and to connect index smoothing to implementability.

5.3. Robust Stability Of Index Anchoring (Exact Bound)

Proposition 1 (Robust stability of index anchoring). *Suppose the closed-loop error dynamics satisfy (9) with bounded disturbance $|w_t| \leq \bar{w} < \infty$. If $0 < \phi < 2$, then the regime is robustly sound in the sense of (6). In particular,*

$$|e_t| \leq |1 - \phi|^t |e_0| + \frac{1 - |1 - \phi|^t}{1 - |1 - \phi|} \bar{w} \leq |1 - \phi|^t |e_0| + \frac{\bar{w}}{1 - |1 - \phi|}. \quad (10)$$

Proof. Iterate (9): $e_t = (1 - \phi)^t e_0 + \sum_{k=0}^{t-1} (1 - \phi)^{t-1-k} w_k$. If $0 < \phi < 2$ then $|1 - \phi| < 1$; apply absolute values and bound the geometric series.³ $\square$

Remark 3 (Bounded disturbances as a design envelope). *The bounded-disturbance assumption is an engineering-style “design envelope,” not a claim that shocks are literally bounded in all states of the world. In applications, $\bar{w}$ (and analogously $\bar{\xi}$) can be set to a high quantile of historical or scenario-based residuals (stress-test magnitudes), yielding verifiable worst-case guarantees within the prescribed envelope. Rare tail events that exceed the envelope may push the system outside the guaranteed bounds; the constitutional design problem then includes specifying preannounced emergency protocols (and corresponding wedge caps) for such exceedances, rather than relying on open-ended discretion.*

³ This is the standard solution/bound for a stable scalar linear difference equation with bounded input; see, e.g., Zhou et al. [20, Ch. 3].

5.4. Choosing the Feedback Gain ϕ : Speed–Noise Trade-Off and Calibration

Equation (9) implies that, absent disturbances ($w_t \equiv 0$), deviations decay at the geometric rate $|1 - \phi|$. A convenient summary is the *half-life* (in model periods) of a deviation:

$$h(\phi) := \frac{\log(1/2)}{\log|1 - \phi|}. \quad (11)$$

For $\phi \in (0, 1)$ the adjustment is monotone (no sign reversals), while $\phi \in (1, 2)$ remains stable but implies alternating overshoots.

In the presence of bounded disturbances, Proposition 1 yields the steady-state sensitivity

$$\limsup_{t \rightarrow \infty} |e_t| \leq \frac{\bar{w}}{1 - |1 - \phi|} = \begin{cases} \bar{w}/\phi, & \phi \in (0, 1], \\ \bar{w}/(2 - \phi), & \phi \in [1, 2). \end{cases} \quad (12)$$

Thus larger ϕ reduces the amplification of persistent shocks *holding $\bar{w}$ fixed*. In practice, however, the effective disturbance term includes measurement noise in e_t and transmission noise in the mapping from the instrument to relative inflation, so aggressive gains can increase instrument volatility and feed high-frequency noise back into w_t .

Calibration recipe (data-based and implementable).

A practical approach is to pick ϕ jointly with the index-smoothing choice in Section 3:

1. **Choose frequency and a target half-life.** Decide whether the regime is evaluated monthly or quarterly and choose a desired mean-reversion speed (e.g., a half-life on the order of several quarters for macro-level stability). Equation (11) provides the implied ϕ ; for instance, a quarterly half-life of $h = 4$ implies $\phi \approx 1 - 2^{-1/4} \approx 0.16$ (monotone adjustment if $\phi < 1$).
2. **Estimate disturbance magnitudes from historical data.** Using historical series for P_t and the constructed anchor I_t , compute e_t and form residuals from the error-correction identity (8): $w_t := (\pi_{t+1} - \pi_{t+1}^I) - \pi^* + \phi e_t$. A conservative bound $\bar{w}$ can be set to a high quantile (e.g., 95th percentile) of $|w_t|$ rather than the maximum, to avoid outlier sensitivity.
3. **Check tolerances and instrument volatility.** Verify that the implied bound in (12) is consistent with an acceptable tolerance for $|e_t|$, and confirm that the implied instrument variability under (15) remains operationally plausible. If rate volatility is excessive, reduce ϕ or increase smoothing in I_t so that measurement noise entering e_t is smaller.

5.5. Mapping (8) to a Policy-Rate Feedback Rule

A minimal implementable mapping is:

$$(\pi_{t+1} - \pi_{t+1}^I) - \pi^* = -\chi(i_t - \bar{i}_t) + u_{t+1}, \quad \chi > 0, \quad (13)$$

where i_t is the policy rate, $\bar{i}_t$ is a (possibly time-varying) neutral nominal rate, and u_{t+1} is bounded.

Rationale. This reduced-form mapping expresses the implementable link from the policy rate to relative inflation (prices in domestic units versus the anchor). The parameter χ summarizes the slope of transmission through aggregate demand and pricing; $\bar{i}_t$ is the neutral nominal rate consistent with the standard absent transitory shocks. The disturbance u_{t+1} captures bounded misspecification and unforeseen shocks. In this form, the feedback rule in (15) delivers the error-correction condition (8) without committing to a full structural model. At the effective lower bound, the rate channel may weaken (small effective χ) or become constrained; in that case the same framework can be extended by allowing balance-sheet interventions to enter as an auxiliary wedge τ_t (Section 9), with explicit bounds and penalties to reflect their distortionary character.

Effective lower bound (ELB) and auxiliary balance-sheet transmission.

In low-rate environments, the policy rate may be constrained ($i_t \geq i^{ELB}$) and the marginal transmission from $(i_t - \bar{i}_t)$ to relative inflation can weaken. A transparent extension is to augment the reduced-form mapping (13) with a balance-sheet wedge τ_t that shifts longer yields and credit conditions:

$$(\pi_{t+1} - \pi_{t+1}^I) - \pi^* = -\chi_i(i_t - \bar{i}_t) - \chi_\tau \tau_t + u_{t+1}, \quad \chi_i, \chi_\tau > 0, \quad (14)$$

where τ_t is the intervention wedge analyzed in Section 9. When the ELB binds, τ_t can close the remaining error-correction gap, but it is explicitly penalized in (21) and constitutionally bounded in (46). Equivalently, define the effective stance $s_t := \chi_i(i_t - \bar{i}_t) + \chi_\tau \tau_t$, so the error-correction condition can be implemented by choosing $s_t = \phi e_t$ subject to the instrument constraints.

Then the feedback rule

$$i_t = \bar{i}_t + \frac{\phi}{\chi} e_t \quad (15)$$

implies (8) with $w_t = u_{t+1}$.

Rationale. This rule translates the abstract error-correction condition into an observable, implementable instrument setting. The rate moves proportionally to the level gap e_t , so policy responds to accumulated past misses rather than only next-period inflation. The neutral-rate term $\bar{i}_t$ separates long-run stance from stabilization; specifying how $\bar{i}_t$ is estimated is part of the constitution, because discretionary redefinition can mimic policy drift. In practice, a transparent starting point is to anchor $\bar{i}_t$ to published neutral-rate (r^*) estimates along the lines of Laubach and Williams [21], Holston et al. [22]. The scaling ϕ/χ makes the closed-loop gain transparent, allowing robustness and tolerance bands to be calibrated.

Remark 4 (Noise, smoothing, and Nash's emphasis on index quality). *A higher gain ϕ corrects deviations faster but can amplify noise in w_t . Index smoothing and refined-index design (Section 3) reduce the effective disturbance magnitude by filtering short-run commodity volatility, consistent with Nash's emphasis on careful index construction and smoothing [10].*

6. A Minimal Monetary Environment And Long-Horizon Contracting Risk

To connect index anchoring to incentives, we use a deliberately minimal reduced-form environment. The purpose is to keep the analysis assumption-light: rather than committing to a fully microfounded New Keynesian structure, we introduce only the minimal ingredients needed to (i) link the policy instrument to the price level, (ii) represent long-horizon contracting as exposure to nominal-value risk, and (iii) represent financial-sector moral hazard as sensitivity to expected discretionary relief.

6.1. Nominal Side (Reduced Form)

We take (13) as a minimal reduced-form mapping from the implementable policy stance to relative inflation (inflation measured against the published anchor). The purpose is not to commit to a particular structural model, but to isolate the closed-loop behavior of the index-anchored standard under bounded misspecification and shocks.

Appendix B provides an alternative operational representation in which the same error-correction logic is implemented through a policy-controlled nominal liability ("base control") under a standard money-demand relation.

Rationale. The subsequent robustness results depend on the induced dynamics for e_t , not on the specific microfoundation of the transmission channel. Presenting a reduced form in the main text keeps the framework auditable and regime-focused, while the appendix documents one conventional implementation mapping for readers who prefer an explicit monetary-liability channel.

6.2. Long-Horizon Contracting Risk: A Horizon- H Statistic

Nash's contracting argument concerns the *future price level* over contract maturities [10]. We therefore replace one-step inflation uncertainty with a horizon-relevant measure.

For horizon $H \geq 1$, define cumulative *relative* inflation between t and $t + H$:

$$\Pi_{t,H} := (\log P_{t+H} - \log P_t) - (\log I_{t+H} - \log I_t) - \pi^* H = \sum_{j=1}^H [(\pi_{t+j} - \pi_{t+j}^I) - \pi^*]. \quad (16)$$

Define the contracting-risk proxy

$$\sigma_{t,H}^2 := \text{Var}_t(\Pi_{t,H}). \quad (17)$$

Definition 5 (Horizon- H contracting-risk statistic). For horizon $H \geq 1$, define cumulative relative inflation $\Pi_{t,H}$ as in (16). Define the contracting-risk statistic $\sigma_{t,H}^2 := \text{Var}_t(\Pi_{t,H})$ as in (17).

Choosing the horizon H (contract maturity as a design parameter).

In applications, H should be tied to the maturities over which the unit of account is economically decisive. Let the model period be Δ years (e.g., $\Delta = 1/4$ for quarterly, $\Delta = 1/12$ for monthly). If the representative contracting horizon is T years, set $H := \lfloor T/\Delta \rfloor$. A practical constitution can either (i) publish a single benchmark horizon H (e.g., a mortgage-dominant horizon), or (ii) report a small *term structure* $\{\sigma_{t,H}^2\}_{H \in \mathcal{H}}$ for a fixed grid $\mathcal{H}$ spanning major nominal exposures (households, corporates, public debt), or (iii) define a composite risk statistic

$$\tilde{\sigma}_t^2 := \sum_{h \in \mathcal{H}} \omega_h \sigma_{t,h}^2, \quad \omega_h \geq 0, \quad \sum_{h \in \mathcal{H}} \omega_h = 1, \quad (18)$$

with weights ω_h proportional to constitutionally defined nominal exposure at horizon h .

Rationale. Nash's contracting criterion is inherently horizon-dependent: the relevant "instability" is uncertainty about the price level over the maturities that govern credit-market structure. Fixing H (or a horizon grid $\mathcal{H}$) by reference to observable maturity distributions makes the criterion operational and prevents ex post relabeling of what counts as "long run."

Estimating $\sigma_{t,H}^2$ in real time.

The object $\text{Var}_t(\Pi_{t,H})$ can be estimated from a constitutionally specified forecasting class for relative inflation $(\pi_{t+1} - \pi_{t+1}^I) - \pi^*$ (e.g., ARMA, local-level state space, or a rolling forecast-error variance with fixed window and outlier rules). When market-based measures exist, the same object can be cross-checked against derivative- or survey-implied distributions of cumulative inflation. In all cases, the estimator class, window length, and any missing-data rules should be published ex ante as part of the measurement constitution.

Rationale. The horizon-risk term is only governance-relevant if it is measured in a reproducible way. Pre-specifying the estimator and update rules limits discretion in the measurement layer and keeps the contracting-risk statistic auditable.

Back-of-envelope magnitude: a high-inflation episode (U.S. CPI-U, 1970–1982).

To give scale, consider the "Great Inflation" era. Using annual CPI-U index averages, one can approximate (log) inflation π_t and fit a simple AR(1) for illustration:

$$\pi_{t+1} - \bar{\pi} = \rho(\pi_t - \bar{\pi}) + \varepsilon_{t+1}, \quad \mathbb{E}(\varepsilon_{t+1}^2) = \sigma_\varepsilon^2. \quad (19)$$

If relative inflation obeys the same reduced-form persistence, then the conditional variance of the H -horizon sum admits the closed-form approximation

$$\sigma_{t,H}^2 = \text{Var}_t \left(\sum_{j=1}^H \pi_{t+j} \right) \approx \sigma_\varepsilon^2 \sum_{k=1}^H \left(\frac{1 - \rho^{H-k+1}}{1 - \rho} \right)^2. \quad (20)$$

Estimating (19) on U.S. annual CPI-U inflation over 1970–1982 yields $\hat{\rho} \approx 0.54$ and $\hat{\sigma}_\varepsilon \approx 0.024$ (log points), implying the horizon scaling in Table 3.⁴

Table 3. Back-of-envelope contracting-risk magnitude in a high-inflation episode (U.S. CPI-U, 1970–1982)

Horizon H (years)	$\text{sd}(\Pi_{t,H})$ (log points)	Approx. level s.d. (%)
1	0.02	2.4
5	0.09	9.8
10	0.15	16.2
15	0.19	21.0

Notes: “Approx. level s.d.” is computed as $100(\exp(\text{sd}) - 1)$ to translate log dispersion into an order-of-magnitude percentage uncertainty in the price level over the horizon. Author’s calculations from BLS CPI-U annual averages [23].

Interpretation. In a historically unstable nominal regime, long-horizon price-level uncertainty can reach economically large magnitudes (tens of percent over mortgage-relevant horizons), which rationally raises term premia and shortens nominal contracting. This is the empirical sense in which Nash’s criterion is horizon-based: the relevant stability object is the distribution of the future *price level*, not merely one-step inflation.

Appendix C provides a compact bond-pricing microfoundation showing how risk premia for nominal claims rise with the variance of the future price level (or cumulative inflation), making $\sigma_{t,H}^2$ a directly relevant distortion statistic for long-horizon contracting.

Remark 5 (Connection to the tracking error). Using (7), $\Pi_{t,H} = e_{t+H} - e_t$. Thus stabilizing e_t in a robust sense also disciplines long-horizon nominal risk; the horizon-risk term captures the residual uncertainty that remains under bounded disturbances and imperfect control.

7. Minimal Distortion Objective as Policy Optimization Criterion

7.1. From Short-Run Inflation Volatility to Horizon Contracting Risk

The relevant nominal risk for long-horizon contracting is uncertainty about the *future price level* over the maturities at which households and firms write commitments, not merely one-period inflation volatility. For that reason, the distortion criterion penalizes the horizon- H contracting-risk statistic $\sigma_{t,H}^2 = \text{Var}_t(\Pi_{t,H})$ defined in (17), where $\Pi_{t,H}$ is cumulative *relative* inflation over H periods as in (16). *Rationale.* Mortgages, long-dated debt, leases, and wage agreements load on the distribution of P_{t+H} (relative to the anchor), so the regime should be evaluated on the dispersion of cumulative inflation over those horizons. Penalizing $\sigma_{t,H}^2$ aligns the design objective with the maturities that govern term premia and the effective planning horizon in the real economy, thereby matching Nash’s contracting emphasis in an operational way.

7.2. Distortion Functional With Explicit Wedge Penalties

Define an intertemporal distortion loss:

$$\mathcal{D} = \mathbb{E} \sum_{t=0}^{\infty} \beta^t \left[\lambda_e e_t^2 + \lambda_\sigma \sigma_{t,H}^2 + \lambda_\tau \tau_t^2 + \lambda_b b_t^2 + \lambda_m (m_t - m^*)^2 + \lambda_\ell (\ell_t - \ell^*)^2 \right]. \quad (21)$$

⁴ Because $\text{Var}_t(\Pi_{t,H})$ is invariant to constant shifts, subtracting a constant drift π^* does not change the variance; the purpose here is order-of-magnitude calibration.

Rationale. This objective collects the regime's measurable distortions into a single design criterion. The e_t term enforces fidelity to the value standard; $\sigma_{t,H}^2$ targets contracting-relevant nominal risk over horizon H . The wedge penalties (τ_t, b_t) internalize that discretionary relative-price manipulation and rescue capacity change incentives ex ante. The remaining terms penalize persistent compositional tilts and leverage away from sustainable benchmarks. Quadratic form is a transparent local approximation that makes trade-offs explicit and supports calibration of constitutional caps and feedback gains.

Here:

- e_t penalizes drift from the value standard.
- $\sigma_{t,H}^2$ penalizes long-horizon nominal risk that raises contract premia and shortens maturities.
- τ_t is a relative-price / term-premium wedge (Section 9).
- b_t is a pardon/transfer-capacity wedge (Section 8).
- m_t and ℓ_t are allocation and leverage proxies (Sections 9 and 10).

7.3. Choosing the Weights λ : Normalization, Calibration, and Constrained Design

The coefficients in (21) are constitutional primitives: they encode the relative social cost assigned to (i) deviations from the value standard, (ii) long-horizon nominal risk, and (iii) intervention wedges and their downstream allocative and leverage consequences. Because the summands have different natural units, it is useful to separate *normalization* (making the loss dimensionless and comparable) from *normative trade-offs* (across stabilized objectives and bounded discretion).

Normalization by explicit tolerances (dimensionless scaling).

Let $(\bar{e}, \bar{\sigma}_H^2, \bar{\tau}, \bar{b}, \bar{m}, \bar{\ell})$ denote constitutional reference magnitudes (e.g., tolerance bands or stress-test bounds) for the corresponding objects. A convenient parameterization is

$$\lambda_e = \frac{\omega_e}{\bar{e}^2}, \quad \lambda_\sigma = \frac{\omega_\sigma}{\bar{\sigma}_H^2}, \quad \lambda_\tau = \frac{\omega_\tau}{\bar{\tau}^2}, \quad \lambda_b = \frac{\omega_b}{\bar{b}^2}, \quad \lambda_m = \frac{\omega_m}{\bar{m}^2}, \quad \lambda_\ell = \frac{\omega_\ell}{\bar{\ell}^2}, \quad (22)$$

so that ω -weights are dimensionless and the loss can be read as a weighted sum of squared “fraction-of-tolerance” deviations (and a variance ratio for $\sigma_{t,H}^2$).

Rationale. Without scaling, comparing penalties across terms can be arbitrary because the objects live in different units (log gaps, variances, yield wedges, shares, leverage proxies). Anchoring the λ 's to explicit tolerances makes the criterion operational: a unit contribution to the loss corresponds to breaching a published band. The remaining degrees of freedom are the dimensionless ω 's, which transparently encode normative priority weights.

Weights as shadow prices in a constrained design problem.

An equivalent interpretation is that (21) is a Lagrangian relaxation of a constrained monetary-constitution problem: pick rule parameters and admissible intervention capacity to minimize expected wedge distortions subject to verifiable stability and resilience requirements. For example, one can view the design as choosing constitutional parameters (index smoothing, feedback gain, and wedge caps) to

$$\min \mathbb{E} \sum_{t=0}^{\infty} \beta^t \left[\omega_\tau (\tau_t / \bar{\tau})^2 + \omega_b (b_t / \bar{b})^2 + \omega_m ((m_t - m^*) / \bar{m})^2 + \omega_\ell ((\ell_t - \ell^*) / \bar{\ell})^2 \right] \quad (23)$$

subject to robust anchoring and resilience constraints (e.g., $0 < \phi < 2$ in Proposition 1 and the stress-test inequality in Theorem 1 with wedge bounds (46)). In this view, the implied multipliers on the performance constraints correspond to the λ 's in (21).

Rationale. The constitution already imposes hard bounds (admissible sets) on wedges and on safe-set excursions; the remaining design problem is to minimize the typical (not worst-case) use of distortionary tools while still meeting robust anchoring and resilience. Treating λ 's as shadow prices makes the trade-offs explicit and connects the loss functional directly to the verifiable inequalities used elsewhere in the paper.

Empirical calibration and reporting discipline.

A practical calibration begins by fixing tolerances $(\bar{e}, \bar{\sigma}_H^2)$ from historical distributions or market-implied measures of long-horizon nominal risk, and by fixing admissible wedge caps $(\bar{\tau}, \bar{b})$ from governance choices (legal authority, balance-sheet capacity, and acceptable incentive effects). The remaining dimensionless weights ω can then be chosen by one of three transparent methods: (i) *moment matching* to empirical relationships (e.g., sensitivity of long yields/term premia and maturity structure to inflation uncertainty; sensitivity of leverage/spreads to wedge episodes); (ii) *stress-test optimization* that selects ω to minimize predicted crisis-frequency proxies subject to the invariance inequalities; or (iii) *welfare mapping* when a structural environment is specified, using a quadratic approximation to social welfare to pin down relative weights. Because the mapping from wedges to outcomes is uncertain, a useful discipline is to report sensitivity of key results to a small grid of ω 's and to display the implied Pareto frontier between contracting stability $(e_t, \sigma_{t,H}^2)$ and wedge-induced distortions $(\tau_t, b_t, m_t, \ell_t)$.

7.4. Incentive Wedges (Retained Conceptual Object)

To formalize “distorted incentives,” we use the notion of an incentive wedge: the gap between the private marginal gain from an action and its social marginal gain. Let a_t denote a choice variable (e.g., leverage, risk, or sectoral allocation). Write the private objective as $U_t(a_t)$ and social welfare as $W_t(a_t)$. The incentive wedge is

$$\Delta_t(a) := \frac{\partial U_t}{\partial a}(a_t) - \frac{\partial W_t}{\partial a}(a_t). \quad (24)$$

Discretionary pardons and targeted wedges enlarge Δ_t by increasing the private return to risk-taking and sectoral concentration without commensurate social benefit.

In v3, $\mathcal{D}$ penalizes observables through which wedges manifest: the wedge processes (τ_t, b_t) and the induced state variables (m_t, ℓ_t, s_t) .

8. The “Pardoner” Channel: Bailout Capacity Wedge And Moral Hazard

Nash explicitly identifies central banks and sovereign issuers as potential “pardoners” of overindebtedness [10]. We represent pardon capacity by a wedge $b_t \geq 0$ that scales the expected private benefit of tail-risk exposure.

8.1. Corrected Stylized Moral-Hazard Model

At time t , an intermediary selects a scalar risk exposure $r_t \in [0, \bar{r}]$ (interpretable as leverage, tail exposure, or risky asset share). Let:

- $\Pi(r)$ be private gross expected returns (increasing, concave),
- $C(r)$ be private costs (increasing, convex),
- $q(r) \in (0, 1)$ be tail-event probability (increasing in r),
- $T(r) \geq 0$ be the transfer/relief conditional on tail events (weakly increasing in r),
- $b_t \geq 0$ be institutional pardon capacity (generosity/feasibility of rescue).

Define expected private payoff:

$$U(r_t; b_t) = \Pi(r_t) - C(r_t) + q(r_t) b_t T(r_t). \quad (25)$$

Proposition 2 (More pardon capacity increases equilibrium risk). *Assume $\Pi - C$ is strictly concave on $[0, \bar{r}]$ and $q'(r)T(r) + q(r)T'(r) \geq 0$. Then any interior optimum r_t^* satisfies the FOC:*

$$\Pi'(r_t^*) - C'(r_t^*) + b_t \left(q'(r_t^*)T(r_t^*) + q(r_t^*)T'(r_t^*) \right) = 0, \quad (26)$$

and r_t^* is weakly increasing in b_t (strictly if the bailout-marginal term is strictly positive at the optimum).

Rationale. The first two terms are the marginal net private return to risk absent rescues. The final term is the marginal expected transfer in tail states: raising risk increases either the probability of distress $q(r)$ or the size of relief $T(r)$, and b_t scales how much of that tail loss is socialized. Because b_t enters the marginal condition, tighter constitutional limits on b_t directly reduce equilibrium risk-taking and weaken the ex ante incentive to load tail exposure [10].

Remark 6 (Institutional implication). *An “ideal money” constitution must bound the pardon-capacity wedge b_t to reduce equilibrium tail-risk selection, consistent with Nash’s critique of pardoners [10]. Emergency facilities can exist, but must be pre-specified, penal, collateralized, and quantitatively capped so that b_t is bounded and priced.*

9. Intervention Wedges And Stage A Distortion Accounting

Two-stage treatment: exogenous distortion accounting vs. endogenous political economy. We separate (i) quantifying how interventions distort incentives from (ii) explaining why interventions occur. In Stage A we treat intervention intensity as exogenous—a bounded sequence chosen outside the model—and derive its mechanical effects on prices, leverage, spreads, and resource allocation. In Stage B we endogenize interventions by introducing a policymaker who chooses them to minimize a short-horizon loss under crisis pressure; constitutional design enters as constraints that reduce the equilibrium frequency and magnitude of interventions.

9.1. A QE-Style Relative-Price Wedge As A Term-Premium / Segmentation Wedge

To avoid an untestable assertion about “market-clearing suppression,” we define τ_t as an explicit wedge in a targeted yield relative to a counterfactual benchmark. Let y_t^H be a targeted yield (e.g., long-duration MBS) and y_t^{cf} a counterfactual yield absent targeted purchases under the same macro state. Define:

$$y_t^H = y_t^{cf} - \tau_t + \varepsilon_t^H, \quad \tau_t \geq 0. \quad (27)$$

Rationale. This equation defines τ_t as the gap between the observed targeted yield and a counterfactual yield absent targeted purchases, holding the macro state fixed. Framed this way, τ_t is a pure relative-price wedge (term-premium/segmentation wedge), not a claim about “market clearing.” It is therefore suitable for distortion accounting: given an estimate of y_t^{cf} , one can quantify how much policy shifted borrowing costs and trace the induced collateral and allocation responses. The residual ε_t^H isolates measurement and model error.

Identification note: y_t^{cf} is not observed; empirically τ_t must be estimated via term-premium models, event studies, or instruments. In this paper τ_t is a wedge for distortion accounting and for constitutional design. Representative term-premium decompositions and LSAP/QE event-study approaches include Kim and Wright [24], Adrian et al. [25], Gagnon et al. [26], Krishnamurthy and Vissing-Jorgensen [27].

9.2. Propagation Block: Demand, Collateral, Leverage, Spreads

Stage A requires a transparent mapping from an intervention wedge to macro-financial outcomes that are directly interpretable as distortions to incentives and allocation. We therefore adopt a deliberately parsimonious propagation block linking the relative-price wedge τ_t to (i) interest-sensitive demand, (ii) collateral price dynamics, (iii) leverage/maturity mismatch, and (iv) credit stress/spreads. The block is designed to satisfy economically standard sign restrictions and to keep each step measurable.

Rationale. The objective is distortion accounting, not full structural identification. A bounded wedge path should imply a bounded and auditable sequence of movements in collateral, leverage, and spreads through a minimal set of channels. This keeps the empirical interpretation clear: τ_t is a relative-price wedge that propagates mechanically into balance-sheet and allocation variables that enter the distortion functional (21) and the resilience constraints in Section 10.

Let D_t denote interest-sensitive demand (e.g., for housing finance), p_t a collateral price index, ℓ_t a leverage/mismatch proxy, and s_t a credit-spread/stress proxy:

$$D_t = \bar{D} - d_y y_t^H + d_z Z_t + v_t^D, \quad d_y > 0, \quad (28)$$

$$\Delta \log p_{t+1} = \alpha(D_t - \bar{D}) + v_{t+1}^p, \quad \alpha > 0, \quad (29)$$

$$\ell_{t+1} = \rho_\ell \ell_t + \beta_p \Delta \log p_{t+1} + \beta_\tau \tau_t + v_{t+1}^\ell, \quad 0 < \rho_\ell < 1, \quad (30)$$

$$s_{t+1} = \rho_s s_t + \gamma_\ell \ell_{t+1} + \gamma_p (\Delta \log p_{t+1})^2 + v_{t+1}^s, \quad 0 < \rho_s < 1. \quad (31)$$

Rationale. This block is a deliberately minimal transmission chain for Stage A distortion accounting. A relative-price wedge in a targeted yield shifts interest-sensitive demand, which moves collateral values; collateral moves relax constraints and expand leverage; leverage and realized price volatility raise stress/spreads. The objective is not a full structural identification of each link, but a parsimonious mapping from the wedge τ_t into observable macro-financial outcomes (p_t, ℓ_t, s_t) under transparent sign restrictions. This makes the mechanical channels and their incentive relevance auditable, and it supports bounding and penalizing wedges in the design criterion.

The sign restrictions formalize the idea that targeted yield wedges raise sector demand, which raises collateral values, which expands leverage, which widens stress/spreads (especially when price growth becomes volatile).

9.3. Bounded Allocation Shares: Logistic Dynamics

Sectoral allocation variables (capital shares, credit shares, or labor-deployment shares) are inherently bounded objects: by definition they lie in $(0, 1)$. To keep the allocation statistic economically meaningful under shocks and wedge episodes, we model the targeted-sector share m_t as a logistic transformation of an unconstrained latent index $z_t \in \mathbb{R}$, so that boundedness holds by construction:

$$m_t = \frac{1}{1 + e^{-z_t}}.$$

Rationale. Using a bounded share avoids artificial behavior (e.g., negative shares or shares exceeding one) and makes deviations comparable across time and jurisdictions. The latent-index dynamics for z_t retain tractability—allowing a simple linear reduced form for how wedges, collateral growth, and stress tilt allocation—while the logistic map guarantees that the measured allocation outcome remains interpretable and directly usable in the penalty term $(m_t - m^*)^2$ in (21).

Let Inv_t^T denote investment (or a resource flow) allocated to the targeted sector and Inv_t^N the corresponding flow allocated elsewhere. Define the sectoral share:

$$m_t := \frac{\text{Inv}_t^T}{\text{Inv}_t^T + \text{Inv}_t^N} \in (0, 1). \quad (32)$$

Rationale. Expressing allocation as a share isolates compositional tilts from aggregate growth and makes deviations comparable across time and jurisdictions. It also matches measurement: national accounts and flow-of-funds data naturally report sector shares of investment, credit, or employment. Because intervention wedges operate through relative prices, they primarily shift shares rather than levels. Using m_t as the allocation statistic therefore links the wedge processes to a directly interpretable misallocation outcome and makes the penalty term $(m_t - m^*)^2$ in (21) economically transparent.

To enforce the $(0, 1)$ bound while preserving tractable dynamics, represent m_t through a latent index $z_t \in \mathbb{R}$ via the logistic map:

$$m_t = \frac{1}{1 + e^{-z_t}}. \quad (33)$$

Let

$$z_{t+1} = \rho_m z_t + \theta_\tau \tau_t + \theta_p \Delta \log p_{t+1} - \theta_s s_t + v_{t+1}^m, \quad 0 < \rho_m < 1. \quad (34)$$

Rationale. Sectoral allocation variables are shares and must remain in $(0, 1)$ to be economically meaningful. Modeling the share through a latent index $z_t \in \mathbb{R}$ with a logistic map guarantees boundedness without ad hoc clipping and preserves smooth comparative statics. The linear dynamics for z_t then provide a tractable reduced-form representation of how wedges, collateral movements, and stress tilt allocation over time, while still allowing mean reversion via $\rho_m < 1$. This construction is particularly convenient when allocation enters the distortion functional as a deviation from a benchmark share.

Analogous dynamics can be used for labor deployment shares.

9.4. Stage A: Distortion Accounting Bound (Schematic)

Given bounded wedge paths $\{\tau_t, b_t\}$ and bounded shocks, (28)–(34) map wedges to trajectories for collateral prices, leverage, spreads, and allocation. The augmented distortion criterion (21) can then be evaluated (or estimated empirically) as a function of wedge magnitudes and persistence.

10. Macro-Financial Resilience: Safe-Set Invariance And Explicit Bounds

Bridge to system-level resilience. Proposition 1 establishes robust stability of the index-tracking error e_t under bounded disturbances, which captures the core “value standard” aspect of ideal money. However, even when e_t remains small, leverage and funding-stress dynamics can generate nonlinear amplification, regime shifts, and prolonged recoveries that shorten contracting horizons and elevate risk premia. To make “soundness” equivalent to shock absorption without tipping, we extend the analysis from a scalar tracking-error recursion to a coupled multi-sector state system and require safe-set invariance and resilience.

10.1. State Vector And Safe Operating Set

Define a compact macro-financial state:

$$x_t := (e_t, \sigma_t, \ell_t, s_t)^\top \in \mathbb{R}^4, \quad (35)$$

where σ_t is a scalar proxy for nominal uncertainty (e.g., a filtered estimator related to $\sigma_{t,H}^2$), ℓ_t is leverage/mismatch, and s_t is stress/spreads.

We define a safe operating set as the region in which long-horizon contracting remains viable (bounded nominal uncertainty and spreads) and in which risk selection is not dominated by runaway leverage feedback:

$$S := \{x : |e| \leq \bar{e}, 0 \leq \sigma \leq \bar{\sigma}, 0 \leq \ell \leq \bar{\ell}, 0 \leq s \leq \bar{s}\}. \quad (36)$$

10.2. Abstract Safe-Set Invariance

We model coupled dynamics abstractly as

$$x_{t+1} = F(x_t) + \xi_t, \quad \|\xi_t\| \leq \bar{\xi}, \quad (37)$$

and impose an invariance margin condition

$$F(S) \oplus B_{\bar{\xi}} \subseteq S. \quad (38)$$

Proposition 3 (Safe-set invariance and resilience). *Under (38), if $x_0 \in S$ then $x_t \in S$ for all $t \geq 0$. Moreover, if F is locally contractive on S (there exists $0 < \rho < 1$ such that $\|F(x) - F(y)\| \leq \rho\|x - y\|$ for all $x, y \in S$), then there exist $C < \infty$ and $0 < \rho < 1$ for which*

$$\|x_t - x^*\| \leq C\rho^t\|x_0 - x^*\| + \frac{C}{1-\rho}\bar{\xi}, \quad (39)$$

for any fixed point $x^* \in S$ of the disturbance-free dynamics.

10.3. Verifiable Sufficient Conditions For Invariance (Linear Or Linearized Case)

The abstract condition (38) is correct but not directly checkable without specifying F . We therefore provide a verifiable sufficient condition under a linear (or linearized) closed-loop state-space model.

Assume

$$x_{t+1} = Ax_t + Bu_t + G\zeta_t, \quad (40)$$

where $u_t := (\tau_t, b_t)^\top$ are policy wedges and ζ_t are exogenous shocks. Suppose

$$\|u_t\|_\infty \leq \bar{u}, \quad \|\zeta_t\|_\infty \leq \bar{\zeta} \quad \forall t. \quad (41)$$

Theorem 1 (Sufficient condition for forward invariance of an ∞ -norm safe box). *Let $\|\cdot\|_\infty$ be the induced matrix norm and define $a := \|A\|_\infty$ and $M := \|B\|_\infty\bar{u} + \|G\|_\infty\bar{\zeta}$. If $a < 1$ and $\bar{x}$ satisfies*

$$M \leq (1 - a)\bar{x}, \quad (42)$$

then $S_\infty(\bar{x}) := \{x : \|x\|_\infty \leq \bar{x}\}$ is forward invariant.

Rationale. Safe-set invariance is the system-level analogue of anchoring: it requires that leverage, stress, and nominal-risk proxies remain within a region compatible with long-horizon contracting, even when shocks and bounded emergency actions occur. The linear sufficient condition turns this into a checkable inequality linking amplification ($\|A\|_\infty$), permitted intervention capacity ($\bar{u}$), and stress-test shock magnitudes ($\bar{\zeta}$). It therefore provides an operational rule for selecting constitutional caps and safe thresholds.

Proof. If $\|x_t\|_\infty \leq \bar{x}$, then $\|x_{t+1}\|_\infty \leq \|A\|_\infty\|x_t\|_\infty + \|B\|_\infty\bar{u} + \|G\|_\infty\bar{\zeta} \leq a\bar{x} + M \leq \bar{x}$. $\square$

10.4. Explicit Trajectory Bound (The “Improvement F ” Bound)

Proposition 4 (Explicit trajectory bound / ISS inequality). *Under (40)–(41) with $a = \|A\|_\infty < 1$ and $M = \|B\|_\infty\bar{u} + \|G\|_\infty\bar{\zeta}$, for all $t \geq 0$,*

$$\|x_t\|_\infty \leq a^t\|x_0\|_\infty + \frac{1 - a^t}{1 - a}M, \quad (43)$$

and in particular

$$\sup_{t \geq 0} \|x_t\|_\infty \leq \|x_0\|_\infty + \frac{M}{1 - a}. \quad (44)$$

Rationale. While invariance gives a yes/no guarantee, the ISS bound quantifies how far the state can move inside the safe region. It decomposes excursions into a decaying component from initial conditions and a steady-state component proportional to the maximum combined input M . This is the natural stress-test object: it maps wedge caps and shock bounds into a numerical upper bound on $\sup_{t \geq 0} \|x_t\|_\infty$, allowing policymakers to choose $\bar{u}$ and $\bar{x}$ to meet explicit resilience tolerances.

Proof. Take norms: $\|x_{t+1}\|_\infty \leq a\|x_t\|_\infty + M$, then iterate and sum the geometric series. $\square$

Remark 7 (Constitutional design as a stress-test inequality). *The bound (44) provides a transparent “stress test” mapping from amplification (a), wedge bounds ($\bar{u}$), and shock bounds ($\bar{\zeta}$) into an upper bound on macro-financial excursions. It formalizes the idea that an ideal-money constitution must both (i) stabilize the unit of account and (ii) bound discretionary wedge capacity to preserve resilience.*

11. Stage B: Endogenous Wedges, Time Inconsistency, and Constitutional Bounds

Stage A quantifies distortions conditional on wedge paths. Stage B explains why wedges occur and how constitutional constraints alter their equilibrium distribution. This is the standard dynamic-

commitment (rules-versus-discretion) problem, here expressed for intervention wedges rather than only inflation bias [7,8].

Takeaway. Stage B does not assume discretion disappears; it assumes crisis pressure makes wedges tempting ex post. A monetary constitution is a commitment device that restricts the admissible intervention set, thereby truncating the realized distribution of wedges and reducing their equilibrium frequency and magnitude.

11.1. A Minimal Time-Inconsistency Sketch

Let a policymaker choose wedges $u_t = (\tau_t, b_t)$ to minimize a short-horizon loss

$$L_t^{pol} = \omega_s s_t^2 + \omega_y (y_t - \bar{y})^2 + \omega_u \|u_t\|^2, \quad (45)$$

subject to feasibility and political constraints. In crises, the marginal benefit of reducing s_t rises, which increases the temptation to choose larger wedges. Yet larger wedges increase future moral hazard (via b_t) and distort allocation (via τ_t), raising $\mathcal{D}$.

11.2. Constitutional Bounds

A monetary constitution imposes admissible-set constraints such as:

$$|\tau_t| \leq \bar{\tau}, \quad 0 \leq b_t \leq \bar{b}, \quad (46)$$

and requires emergency facilities to be penal, collateralized, and rule-triggered, so that wedge realizations are bounded and predictable in form.

Remark 8. In this framework, constitutional bounds do not eliminate crisis response; they constrain its scale and pricing, shifting the equilibrium distribution of wedges toward smaller and less persistent interventions, thereby reducing both the moral-hazard term (Proposition 2) and the allocation-tilt dynamics (Section 9).

12. Implementation and Measurement

12.1. Measuring the Value-Standard Gap

The tracking error e_t is directly computable from published P_t and the published anchor I_t :

$$e_t = \log(P_t / I_t) - (\log \kappa + \pi^* t).$$

Practical design choices include: which P_t (headline CPI, trimmed mean, or a chain-weighted deflator), how frequently I_t is published (daily vs monthly), and whether e_t is evaluated with real-time or revised data.

12.2. Measuring Horizon Risk

The key statistic is $\sigma_{t,H}^2 = \text{Var}_t(\Pi_{t,H})$. Empirically, Var_t can be proxied using: (i) inflation-derivative implied measures; (ii) survey dispersion over cumulative inflation; or (iii) rolling forecast-error variance from an agreed estimator. The constitutional design should specify the estimator class to avoid politicized redefinition.

12.3. Estimating Wedges

The wedge τ_t requires a counterfactual yield estimate; standard approaches include term-premium decompositions and high-frequency event studies. The wedge b_t can be proxied by: (i) market-implied probabilities of support (e.g., CDS basis differences), (ii) stated fiscal backstop authority, or (iii) observed intervention rules and caps. For concrete implementations, see term-premium based counterfactual yield estimation in Kim and Wright [24], Adrian et al. [25] and yield effects of asset purchases in Gagnon et al. [26], Krishnamurthy and Vissing-Jorgensen [27], while market-implied support uplifts and TBTF subsidy measures are developed in, e.g., Ueda and Weder di Mauro [28], Santos [29], Acharya et al. [30].

13. Discussion And Policy Implications

The central claim of this paper is that “sound money” should be formulated as an institutional design problem with explicit mathematical criteria, rather than as an aspiration stated only in qualitative terms. Nash’s core vision is that the monetary unit must function as a reliable yardstick for long-horizon contracting and that a principal failure mode of modern regimes is the capacity for discretionary debasement or ad hoc relief—a capacity that acts as an implicit “pardon” for overindebtedness and, crucially, is anticipated ex ante [10]. In that setting, distortions are not limited to a change in the average inflation rate. They arise because agents adapt to the rules of the game: when the unit of account is unreliable or the institution is expected to intervene in sector-specific or bailout-like ways, private choices shift toward shorter horizons, defensive hedging, leverage, speculative positioning, and political rent-seeking. The methodological objective here is to treat those incentive distortions as measurable objects and to incorporate them directly into the definition of what an “ideal” monetary regime should accomplish.

A nonpolitical value standard requires both a formula and a constitution.

A value standard is “nonpolitical” only if it is mechanically computable from transparent data and if the rules governing its evolution cannot be altered ad hoc. This motivates treating the anchor as a published index I_t built from internationally traded commodity prices, accompanied by explicit rules for smoothing and for scheduled, rule-based revisions. The smoothing and “refined index” ideas are not mere statistical preferences; they are institutional safeguards. A regime that reacts to a noisy anchor invites either excessive instrument volatility or selective, discretionary deviations justified as “temporary.” Conversely, an anchor that is refined by transparent filters reduces transitory volatility while preserving long-run relevance. Nash explicitly emphasizes both ingredients: the anchor should be grounded in internationally referenced industrial values and refined to vary smoothly at high frequency while remaining appropriate over long horizons [10]. A constitutional protocol (eligibility, weights, revision triggers, and publication) is therefore part of monetary design, not a separate administrative detail.

From qualitative soundness to a measurable value standard.

A first step is to operationalize price stability as preservation of the monetary unit’s real value over planning-relevant horizons. This is done by anchoring the domestic price level P_t to an objective benchmark index I_t and defining a tracking error $e_t = \log(P_t/I_t) - (\log \kappa + \pi^* t)$. In this formulation, “soundness” is no longer rhetorical: it becomes a control objective, and the success of the regime can be evaluated by how tightly and how robustly it keeps e_t near zero. The robust-soundness inequality makes this precise: the error at time t is bounded by a decaying component that reflects initial misalignment and by a component proportional to the largest shock the system must absorb. The key conceptual shift is that stability is expressed as a bounded-gain property: the unit of account remains reliable under bounded disturbances, rather than requiring perfect forecasting or continuous discretionary fine-tuning.

This definition also clarifies the role of the drift parameter π^* . Nash notes that a currency can be “too good” as a store of value if it becomes a costless global safe-deposit box; a constant, contractible inflation drift can avoid this without sacrificing predictability for long-horizon contracting [10]. In the present framework, π^* is the design parameter that separates “predictability” from “zero drift”: a regime can be ideal for contracting even when it permits a steady, preannounced drift that can be priced into nominal contracts.

Long-horizon contracting risk is inherently multi-period.

A second implication of Nash’s argument is that the relevant measure of nominal instability is not one-step inflation volatility, but the uncertainty in the price level over contract maturities. A household signing a mortgage or a firm issuing long-dated debt is exposed to the distribution of P_{t+H} , not merely to π_{t+1} . For that reason, the contracting-risk statistic $\sigma_{t,H}^2 = \text{Var}_t(\Pi_{t,H})$, where $\Pi_{t,H}$ is cumulative

relative inflation over horizon H , is the natural object to penalize. This term serves as a bridge between monetary stability and credit-market structure: when horizon risk rises, nominal term premia and risk premia rise, maturities shorten, and contracting becomes more defensive. In this sense, a regime can satisfy an ‘inflation target’ in a narrow one-period sense while still failing Nash’s criterion if it generates large uncertainty about the future price level over planning horizons. The design objective therefore combines level tracking (small e_t) with bounded horizon risk (small $\sigma_{t,H}^2$), reflecting the practical reality that long-horizon contracting is governed by the distribution of cumulative inflation.

Why incentives must be part of the definition of ideal money.

However, stabilizing e_t alone does not fully address Nash’s second concern: the distortive effects of discretionary “pardon capacity.” In practice, large-scale interventions can alter relative prices and reshape private behavior in ways that are not captured by headline inflation statistics. A central example is targeted asset purchases that lower sector-specific yields relative to counterfactual price discovery. When mortgage-related instruments are purchased at scale, the mortgage rate can be written as $r_t^{mort} = r_t^{cf} - \tau_t + \varepsilon_t$, where τ_t is an intervention wedge measuring the degree to which yields are reduced relative to an estimated counterfactual benchmark. This wedge has an incentive interpretation: it increases interest-sensitive demand and borrowing capacity in the targeted sector as if there were additional private demand at each price, thereby amplifying collateral values and encouraging leverage. The consequence is a domino effect through the economy: sectoral credit conditions feed into asset prices, asset prices feed into collateral, collateral feeds into balance-sheet expansion, and balance-sheet expansion feeds into systemic risk and credit spreads.

Historical vignette: informal external indicators under discretion.

Greenspan defined price stability as an environment in which the public need not take expected general price changes into account in decisionmaking [31]; market indicators (including commodity prices) were sometimes discussed as contemporaneous signals of inflation expectations. For example, in congressional testimony, Greenspan noted that movements in the price of gold can be broadly reflective of inflationary expectations.⁵ This is not a claim of an operational gold target; rather, it illustrates the practical appeal of an externally generated nominal reference when credibility is fragile. The difference in the present framework is institutional: the reference is elevated from an informally interpreted signal to a published index anchor with explicit construction and revision governance, and the intervention set is constitutionally bounded.

Modern context: QE and asset-purchase programmes as explicit wedges.

In the language of this paper, large-scale asset purchases—including programmes commonly described as “QE” in the United States (Treasury and agency MBS purchases) and the ECB’s Asset Purchase Programme (APP)—enter primarily as a relative-price wedge τ_t : the degree to which targeted yields are compressed relative to a counterfactual benchmark term premium (Section 9). The point is not that such interventions are never justified in crisis states, but that they are not neutral instruments: a persistent τ_t mechanically propagates through collateral values, leverage, and sectoral allocation channels, and therefore can change ex ante incentives and the political demand for continued support. The constitutional implication is that if asset purchases are permitted as an auxiliary tool (for example when the policy-rate channel is impaired), they should be governed by rule-triggered activation, published maturity/eligibility classes, quantitative caps, and transparent exit conditions—so that the induced wedge path is bounded, auditable, and priced in advance rather than open-ended and discretionary. Evidence that large-scale purchase programmes compress targeted yields and term premia is documented for U.S. LSAPs and euro-area APP-style purchases in Gagnon et al. [26], Krishnamurthy and Vissing-Jorgensen [27], Altavilla et al. [33], Eser and Schwaab [34].

⁵ Greenspan [32].

Backstops and “TBTF” support expectations as the b_t channel.

A complementary modern channel is expected official support for particular intermediaries or sectors, which maps to the pardon-capacity wedge b_t (Section 8). In practice, b_t is proxied by market measures of support expectations—for example, “TBTF” funding premia (funding-cost differentials between systemically important and matched non-systemic institutions), ratings-based support uplift, or event-study shifts around guarantee and facility announcements. The design implication parallels the τ_t case: emergency capacity can exist, but a constitution should make it penal, collateralized, and quantitatively capped so that it enters the system as a bounded disturbance rather than an unpriced subsidy to tail-risk selection. Empirical strategies to quantify TBTF funding subsidies and official-support uplift are provided in, e.g., Ueda and Weder di Mauro [28], Santos [29], Acharya et al. [30].

To avoid loading normative claims into definitions, the wedge is best interpreted as a term-premium/segmentation wedge: it measures the difference between an observed yield and a counterfactual yield absent targeted interventions, conditional on the same macro state. This interpretation has two advantages. First, it is aligned with empirical practice (term-premium estimation and event studies). Second, it makes the constitutional design problem explicit: if a regime permits large and persistent wedges, it will predictably induce sectoral tilts and balance-sheet responses even if the aggregate price level remains close to the anchor. In short, incentive distortions can coexist with level stability; a definition of ideal money must therefore constrain the wedge mechanisms through which discretion operates.

Pardons, moral hazard, and the necessity of bounding rescue capacity.

Similarly, discretionary rescues operate as a pardon/transfer wedge—an increase in the expected probability or generosity of relief in tail states—which changes risk selection ex ante. These mechanisms generate the practical phenomena often described as “impaired price discovery” and “socialised losses,” but here they are expressed as explicit wedges that can be measured, bounded, and penalized.

The key structural requirement is that pardon capacity must enter private marginal incentives. When institutional rescue capacity increases, private decision-makers rationally shift toward higher tail exposure because the private downside is partially absorbed by the state in adverse states. This is precisely the channel Nash highlights: the existence of “grand pardoners” changes the game being played and therefore changes behavior [10]. The implication is not that emergency liquidity should be forbidden; rather, emergency capacity must be rule-triggered, penal, collateralized, and quantitatively capped so that it functions as a bounded disturbance to the system rather than an unpriced, state-contingent subsidy to risk-taking.

Formalizing incentive distortions: wedges, outcomes, and a distortion functional.

The paper’s main methodological contribution is to formalize changes in incentive as a gap between privately optimal choices and socially productive behavior, and to connect that gap to observable variables. In abstract terms, distorted incentives correspond to an incentive wedge: the difference between the private marginal payoff from an action (such as leverage, sectoral concentration, or tail risk exposure) and the social marginal payoff once externalities and distributional costs are internalized. While social welfare is not directly observed, the channels through which wedges manifest are observable: intervention wedges, the probability or generosity of discretionary relief, and the downstream tilts in capital allocation and labor deployment. This motivates an augmented distortion functional that penalizes not only anchor deviation but also contracting-relevant nominal risk and the intervention wedges and compositional shifts through which discretion changes behavior.

Two design choices are crucial here. First, penalizing $\sigma_{t,H}^2$ rather than only one-step inflation volatility aligns the criterion with contracting horizons. Second, penalizing the wedge processes (τ_t, b_t) is a constitutional commitment: it expresses the principle that even “successful” crisis stabilization can be socially costly if it systematically induces leverage and sectoral concentration that raises future fragility. In other words, the objective is not merely to minimize short-run stress conditional on crisis,

but to minimize the *equilibrium frequency and severity* of states in which crises arise and discretionary wedges appear attractive.

Resilience and safe-set invariance: preventing the return of discretion through crisis pressure.

A further implication of Nash's critique is that a regime must be judged by its behavior under stress. Even a well-specified index anchor can be undermined if shocks push the macro-financial system into a region where spreads, leverage, and uncertainty become self-reinforcing and policymakers claim that extraordinary discretion is unavoidable. For this reason the soundness requirement extends from scalar tracking stability to a system-level resilience requirement: the macro-financial state $x_t = (e_t, \sigma_t, \ell_t, s_t)$ should remain within a safe operating set under bounded shocks and bounded emergency capacity. Invariance of the safe set is the formal counterpart of "no tipping": the regime is designed so that plausible disturbances do not push the system into regions that trigger repeated departures from rule-based governance.

Two points sharpen the resilience requirement. First, it admits a verifiable stress-test form under linear (or linearized) dynamics: the safe set can be selected to satisfy explicit inequalities linking amplification (matrix norms), wedge bounds, and shock bounds to maximum excursions. Second, it provides a disciplined institutional interpretation: emergency facilities should be designed so that their maximum scale is known *ex ante* and small enough that safe-set invariance can be maintained under calibrated disturbance magnitudes. In this way, resilience is not an aspiration but an engineering constraint: one chooses admissible wedge bounds and operational triggers so that the coupled macro-financial system remains in the region in which long-horizon contracting remains viable and political pressure for pardons remains limited.

A two-stage view of interventions: accounting first, political economy second.

The framework also clarifies how to incorporate modern interventions without blurring causality. In a first stage, policy wedges are treated as exogenous bounded sequences, allowing clean distortion accounting: for any bounded intervention path, one can trace and quantify the induced deviations in spreads, leverage, and sectoral allocation and the resulting increase in the augmented distortion criterion. This is useful even when one does not model the political or institutional process generating interventions. In a second stage, wedges are endogenized by modeling a policymaker (or coalition) that trades off near-term objectives against longer-run distortions. In that setting, constitutional design appears as constraints on admissible interventions, which alter the equilibrium distribution of wedges. The main point is that rules should change not only outcomes conditional on intervention, but also the frequency and magnitude of intervention in equilibrium.

This perspective aligns tightly with Nash's institutional emphasis: a nonpolitical value standard removes a key channel through which pardons occur (general inflationary debasement), but it does not by itself prevent wedge creation through other instruments. A monetary constitution must therefore specify both the value standard and the admissible set of interventions that can be deployed around it. The design goal is not the elimination of all discretion, but the construction of bounded, predictable, penal forms of crisis response that do not create large unpriced incentives for future risk selection.

Concluding perspective.

An "ideal money" regime in the sense suggested by Nash can be understood as a monetary constitution with three intertwined properties. First, it provides a nonpolitical value standard and a transparent feedback mechanism that keeps the unit of account aligned with that standard in a robust, bounded-gain sense. Second, it minimizes the incentive distortions induced by nominal risk and by moral hazard associated with expected discretionary relief, by explicitly penalizing and constraining the policy wedges through which modern interventions alter relative prices and socialise losses. Third, it is resilient: it keeps the macro-financial system within a safe operating region under plausible shocks so that crises do not repeatedly justify the reintroduction of discretion. Taken together, these elements convert the historical debate over "sound money" from a primarily philosophical dispute

into a program of formal design: specify the value anchor, specify the admissible intervention set, and evaluate regimes by their provable robustness and by their quantified effects on incentives, allocation, and long-horizon contracting. Leveraging the formidable capabilities of contemporary computation, we hope the moment may finally have arrived to subject John Nash’s groundbreaking ideas to rigorous, empirical scrutiny.

14. Conclusions

We believe this paper advances Nash’s “ideal money” program by translating it into a policy-implementable and mathematically verifiable institutional design. The central object is an objective value standard and a transparent feedback mechanism that stabilizes the *unit of account* for long-horizon contracting, while explicitly representing (and constraining) the channels through which discretion generates moral hazard and allocative distortions [10]. The resulting criterion for “ideal money” is not reducible to a single inflation statistic: it is a joint requirement of anchor tracking, bounded long-horizon nominal risk, bounded intervention wedges, and macro-financial resilience under stress.

Summary of design requirements and results. An “ideal money” regime in this framework is characterized by the following elements:

1. **Nonpolitical value standard with governance.** A transparent index anchor I_t constructed from internationally traded prices, with explicit smoothing/refinement and a constitutional revision protocol that minimizes susceptibility to ad hoc political manipulation [10].
2. **Implementable feedback rule with robust guarantees.** A rule-based policy mechanism (e.g., an error-correction policy-rate rule) that yields bounded-gain stability of the tracking error e_t under bounded disturbances, so deviations from the value standard remain bounded and mean-reverting.
3. **Contracting-relevant nominal stability.** A minimal-distortion criterion that penalizes *long-horizon* nominal risk (e.g., $\sigma_{t,H}^2 = \text{Var}_t(\Pi_{t,H})$) rather than focusing only on one-step inflation volatility, aligning the policy objective with the maturities that matter for real investment and credit-market structure.
4. **Explicit intervention wedges and constitutional bounds.** An explicit representation of interventions as wedges—a relative-price/term-premium wedge τ_t and a pardon/transfer-capacity wedge b_t —together with admissible-set bounds that limit moral hazard and sectoral misallocation.
5. **System-level resilience.** A resilience requirement stated as safe-set invariance (and supported by verifiable sufficient conditions and trajectory bounds under linear/linearized dynamics), ensuring that plausible shocks do not repeatedly force departures from rule-based governance.

Implementation and research agenda. Operationalizing the design requires (i) specifying the published index methodology (component eligibility, weights, smoothing filter, and revision schedule), (ii) choosing contract-relevant horizons H and estimators for $\sigma_{t,H}^2$ that are themselves rule-governed, and (iii) calibrating constitutional wedge bounds (e.g., $\bar{\tau}, \bar{b}$) to satisfy the resilience inequalities under stress-test shock magnitudes. Empirically, the framework can be evaluated by whether reductions in e_t variability and $\sigma_{t,H}^2$ are associated with longer maturities, lower nominal risk premia, and less persistent sectoral tilting during intervention episodes. These steps convert Nash’s normative vision into a program that is testable, auditable, and institutionally actionable [10].

Next steps (optional). Empirical calibration of (λ, ϕ, H) and wedge bounds $(\bar{\tau}, \bar{b})$ can be framed as minimizing estimated $\mathcal{D}$ subject to resilience constraints of the form in Theorem 1 and Proposition 4.

Appendix A. Notation Summary (Selected)

Symbol	Meaning
P_t	Domestic price level
I_t	Published smoothed commodity anchor
$\hat{I}_t^*$	Raw international ICPI in reference numeraire
E_t	Exchange rate (domestic per reference)
e_t	Tracking error $\log(P_t/I_t) - (\log \kappa + \pi^* t)$
π_{t+1}	Inflation $\log P_{t+1} - \log P_t$
π_{t+1}^I	Anchor inflation $\log I_{t+1} - \log I_t$
$\Pi_{t,H}$	Cumulative relative inflation over horizon H
$\sigma_{t,H}^2$	$\text{Var}_t(\Pi_{t,H})$ (contracting risk)
τ_t	Relative-price / term-premium wedge
b_t	Pardon/transfer-capacity wedge
ℓ_t	Leverage/mismatch proxy
s_t	Credit stress/spread proxy

Appendix B. Base-Control Implementation (A Monetary-Base Feedback Rule)

The main text expresses implementation using an interest-rate instrument, which is the dominant operating framework in modern inflation-targeting regimes. This appendix provides an alternative but equivalent implementation channel in which policy controls a nominal liability (e.g., the monetary base or an operating target for reserves). Under a standard money-demand relation, a simple base-growth feedback rule induces the same closed-loop error recursion for the index-anchored level gap e_t .

A base-growth feedback rule.

Let M_t denote a policy-controlled nominal liability (interpretable as the monetary base or an operating target). Let $\varphi > 0$ be a feedback gain and let u_t capture bounded implementation error. Define the policy update:

$$\Delta \log M_{t+1} = \Delta \log I_t + \pi^* - \varphi e_t + u_t. \quad (\text{A1})$$

A reduced-form nominal relation.

Assume a standard log money-demand representation:

$$\log P_t = \log M_t - \ell_t + \varepsilon_t, \quad (\text{A2})$$

where ℓ_t is (log) real money demand and ε_t is a nominal disturbance. This reduced form is consistent with many environments in which real money demand depends on output and opportunity cost, while ε_t captures velocity and measurement shocks.

Induced error dynamics.

Taking first differences of (A2) yields

$$\pi_{t+1} \equiv \Delta \log P_{t+1} = \Delta \log M_{t+1} - \Delta \ell_{t+1} + \Delta \varepsilon_{t+1}.$$

Substituting (A1) gives

$$\pi_{t+1} = \Delta \log I_t + \pi^* - \varphi e_t + \underbrace{\left(u_t - \Delta \ell_{t+1} + \Delta \varepsilon_{t+1} \right)}_{\text{bounded shocks}}.$$

Using anchor inflation $\pi_{t+1}^I \equiv \Delta \log I_{t+1}$, we can write the relative-inflation error-correction form

$$(\pi_{t+1} - \pi_{t+1}^I) - \pi^* = -\varphi e_t + w_t,$$

where

$$w_t \equiv u_t - \Delta \ell_{t+1} + \Delta \varepsilon_{t+1} + (\Delta \log I_t - \Delta \log I_{t+1})$$

collects bounded implementation error, money-demand/velocity shocks, and the (bounded) timing discrepancy between $\Delta \log I_t$ used in (A1) and realized anchor inflation $\Delta \log I_{t+1}$.⁶

Combining this error-correction relation with the identity linking the level gap to relative inflation yields the same closed-loop recursion as in the main text:

$$e_{t+1} = (1 - \varphi)e_t + w_t.$$

Thus, in a base-control operating regime, M_t should be interpreted as a policy-controlled nominal liability (base or operating target), not broad money; the robust bounded-gain analysis applies verbatim with gain φ and disturbance sequence w_t defined above.

Appendix C. Bond-Pricing Microfoundation For The Horizon-Risk Term

This appendix supplies a compact justification for using $\sigma_{t,H}^2 = \text{Var}_t(\Pi_{t,H})$ as a contracting-risk statistic.

Let a nominal claim pay 1 unit of currency at $t + H$. Its real payoff is $1/P_{t+H}$. Under a lognormal approximation for $\log P_{t+H}$, the certainty-equivalent real payoff declines with $\text{Var}_t(\log P_{t+H})$, implying a higher required nominal yield/spread for longer maturities when price-level uncertainty rises. Because

$$\log P_{t+H} = \log P_t + \sum_{j=1}^H \pi_{t+j},$$

it follows that term premia for nominal claims rise with $\text{Var}_t(\sum_{j=1}^H \pi_{t+j})$. In the index-anchored regime, the relevant object is cumulative *relative* inflation (16), hence $\sigma_{t,H}^2$.

Appendix D. Refined Indices (Nash's "Pricing Modulo" Concept)

Nash suggests that an index can be constructed to have the short-run smoothness of one basket while following the long-run behavior of another by taking a moving average of their ratio ("pricing modulo" a smoother index) [10].

Let $\hat{I}_t^{(B)}$ be a "base/industrial" basket and $\hat{I}_t^{(S)}$ a smoother basket. Define

$$R_t := \log \hat{I}_t^{(B)} - \log \hat{I}_t^{(S)}.$$

Let $\text{MA}_L(R_t)$ be an L -period moving average. Then a refined anchor can be defined as

$$\log I_t := \log \hat{I}_t^{(S)} + \text{MA}_L(R_t). \quad (\text{A3})$$

This preserves short-run smoothness (from $\hat{I}^{(S)}$) while tracking long-run industrial costs (from $\hat{I}^{(B)}$).

Appendix E. Additional Resilience: Minimal Invariant Box From ISS Bound

From Proposition 4, if $a = \|A\|_\infty < 1$ and $M = \|B\|_\infty \bar{u} + \|G\|_\infty \bar{\xi}$, then any box $S_\infty(\bar{x})$ with $\bar{x} \geq \frac{M}{1-a}$ is forward invariant. This yields a constructive method for selecting safe-set thresholds consistent with wedge bounds and stress-test shock magnitudes.

⁶ If desired, the timing term can be reduced by replacing $\Delta \log I_t$ with a constitutionally specified forecast $\pi_{t+1|t}^I$; the difference $\pi_{t+1}^I - \pi_{t+1|t}^I$ is then absorbed into w_t .

References

1. Phillips, A.W. The Relation Between Unemployment and the Rate of Change of Money Wage Rates in the United Kingdom, 1861–1957. *Economica* **1958**, *25*, 283–299. <https://doi.org/10.1111/j.1468-0335.1958.tb00003.x>.
2. Keynes, J.M. *The General Theory of Employment, Interest and Money*; Macmillan: London, 1936.
3. Samuelson, P.A.; Solow, R.M. Analytical Aspects of Anti-Inflation Policy. *American Economic Review (Papers and Proceedings)* **1960**, *50*, 177–194.
4. Friedman, M. The Role of Monetary Policy. *American Economic Review* **1968**, *58*, 1–17.
5. Lucas, R.E. Econometric Policy Evaluation: A Critique. *Carnegie-Rochester Conference Series on Public Policy* **1976**, *1*, 19–46. [https://doi.org/10.1016/S0167-2231\(76\)80003-6](https://doi.org/10.1016/S0167-2231(76)80003-6).
6. Federal Reserve History. Federal Reserve Reform Act of 1977. <https://www.federalreservehistory.org/essays/fed-reform-act-of-1977>, 2013. Federal Reserve History essay. (Cites the 1977 statutory language amending the Federal Reserve Act.).
7. Kydland, F.E.; Prescott, E.C. Rules Rather than Discretion: The Inconsistency of Optimal Plans. *Journal of Political Economy* **1977**, *85*, 473–491.
8. Barro, R.J.; Gordon, D.B. Rules, Discretion and Reputation in a Model of Monetary Policy. *Journal of Monetary Economics* **1983**, *12*, 101–121.
9. Taylor, J.B. Discretion versus Policy Rules in Practice. *Carnegie-Rochester Conference Series on Public Policy* **1993**, *39*, 195–214.
10. Nash, J.F. Ideal Money. *Southern Economic Journal* **2002**, *69*, 4–11.
11. Clarida, R.; Galí, J.; Gertler, M. Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory. *Quarterly Journal of Economics* **2000**, *115*, 147–180.
12. Hansen, L.P.; Sargent, T.J. *Robustness*; Princeton University Press: Princeton, NJ, 2008.
13. Farhi, E.; Tirole, J. Collective Moral Hazard, Maturity Mismatch, and Systemic Bailouts. *American Economic Review* **2012**, *102*, 60–93. <https://doi.org/10.1257/aer.102.1.60>.
14. Svensson, L.E.O. Price-Level Targeting versus Inflation Targeting: A Free Lunch? *Journal of Money, Credit and Banking* **1999**, *31*, 277–295.
15. Eggertsson, G.B.; Woodford, M. The Zero Bound on Interest Rates and Optimal Monetary Policy. *Brookings Papers on Economic Activity* **2003**, 2003.
16. Woodford, M. *Interest and Prices: Foundations of a Theory of Monetary Policy*; Princeton University Press: Princeton, NJ, 2003.
17. Shelton, J. Statement of Judy Shelton (Nominee to be a Member, Board of Governors of the Federal Reserve System). Hearing before the Committee on Banking, Housing, and Urban Affairs, U.S. Senate, 116th Congress, 2nd Session (Feb. 13, 2020), 2020. GovInfo: CHR-116shrg40240.
18. International Monetary Fund. Primary Commodity Prices. IMF Data, 2026. Accessed 2026-01-18.
19. World Bank. Commodity Markets (“Pink Sheet”) Data. World Bank Data, 2026. Accessed 2026-01-18.
20. Zhou, K.; Doyle, J.C.; Glover, K. *Robust and Optimal Control*; Prentice Hall, 1996.
21. Laubach, T.; Williams, J.C. Measuring the Natural Rate of Interest. *Review of Economics and Statistics* **2003**, *85*, 1063–1070.
22. Holston, K.; Laubach, T.; Williams, J.C. Measuring the Natural Rate of Interest: International Trends and Determinants. *Journal of International Economics* **2017**, *108*, S59–S75. Supplement 1.
23. U.S. Bureau of Labor Statistics. Consumer Price Index for All Urban Consumers (CPI-U): U.S. City Average, All Items (Series CUUR0000SA0). BLS Time Series, 2026. Accessed 2026-01-18.
24. Kim, D.H.; Wright, J.H. An Arbitrage-Free Three-Factor Term Structure Model and the Recent Behavior of Long-Term Yields and Distant-Horizon Forward Rates. Finance and Economics Discussion Series 2005-33, Board of Governors of the Federal Reserve System, 2005.
25. Adrian, T.; Crump, R.K.; Moench, E. Pricing the Term Structure with Linear Regressions. *Journal of Financial Economics* **2013**, *110*, 110–138.
26. Gagnon, J.; Raskin, M.; Remache, J.; Sack, B. The Financial Market Effects of the Federal Reserve’s Large-Scale Asset Purchases. *International Journal of Central Banking* **2011**, *7*, 3–43.
27. Krishnamurthy, A.; Vissing-Jorgensen, A. The Effects of Quantitative Easing on Interest Rates: Channels and Implications for Policy. *Brookings Papers on Economic Activity* **2011**. Fall 2011.
28. Ueda, K.; Weder di Mauro, B. Quantifying Structural Subsidy Values for Systemically Important Financial Institutions. *Journal of Banking & Finance* **2013**, *37*, 3830–3842.

29. Santos, J.A.C. Evidence from the Bond Market on Banks' Too-Big-To-Fail Subsidy. Technical report, Federal Reserve Bank of New York, 2014. Working paper / Economic Policy Review version.
30. Acharya, V.V.; Anginer, D.; Warburton, A.J. The End of Market Discipline? Investor Expectations of Implicit Government Guarantees. Technical report, Working paper, 2016.
31. Greenspan, A. Achieving Price Stability. FRASER, Federal Reserve Bank of St. Louis, 1996. Remarks at the Federal Reserve Bank of Kansas City symposium, *Achieving Price Stability*, August 30, 1996.
32. Greenspan, A. Semiannual Monetary Policy Report to the Congress: Testimony before the Subcommittee on Economic Growth and Credit Formation of the Committee on Banking, Finance and Urban Affairs, U.S. House of Representatives. FRASER, Federal Reserve Bank of St. Louis, 1993. July 20, 1993.
33. Altavilla, C.; Carboni, G.; Motto, R. Asset Purchase Programmes and Financial Markets: Lessons from the Euro Area. ECB Working Paper Series 1864, European Central Bank, 2015.
34. Eser, F.; Schwaab, B. Evaluating the Impact of Unconventional Monetary Policy Measures: Empirical Evidence from the European Central Bank's Securities Markets Programme. *Journal of Financial Economics* **2016**, *119*, 147–167.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.