

From Triangulation to Prediction: A State-Space Model of Inference Under Managed Information

Mathematical Foundations, Empirical Calibration, and the Dated Prediction

J. Kenneth Rhodes

Rhodes Research | rhodesresearch.org

jkr@rhodesresearch.org

September 2026

Overview

This paper provides the formal mathematical architecture underlying the five-channel predictive framework for forecasting Chinese policy outcomes and presents its first empirical calibration against observed data. The framework is expressed as a nonlinear, non-Gaussian state-space model with two latent factors — behavioral reality and the official narrative — observed through five channels with heterogeneous loadings, channel-specific lags, and ordinal measurement. Three results anchor the formalization. First, the five-channel architecture is the minimal measurement system for identifying a two-factor reality-narrative structure: the Ledermann bound requires at least five indicators when both factor loadings and idiosyncratic variances are unknown, and the informal three-channel convergence rule corresponds exactly to the Anderson-Rubin three-indicator condition for single-factor identification. Second, the silence map is formalized as informative (non-ignorable) censoring in which the probability of suppression is increasing in the latent narrative-reality gap. Third, the contradiction index is defined as the filtered latent gap, and prediction is implemented as a competing-risks hazard system over three terminal regimes: policy reversal, escalating control, and delayed recognition. The second part presents the first empirical calibration of the framework, exploiting the March-June 2026 commodity shock as a natural identification instrument. The behavioral-reality factor is statistically motionless across the largest exogenous cost shock since 2021 (range -1.32 to -1.21 , $sd \approx 0.15$), while the narrative factor rises from -0.6 to $+4.6$, and the contradiction index expands from 0.65 to 5.90 — a nine-fold widening in six months, robust across the full pass-through sensitivity grid. Because the narrative's support is borrowed from year-on-year arithmetic, its decay is deterministic under flat commodity levels, dating a forced narrative-correction window in March-June 2027 and gating the competing-risks hazard layer with a calendar rather than a guess. Falsification conditions for the dated prediction are specified as formal interventions on model parameters, mapped to the three terminal regimes.

PART I: THE FORMAL ARCHITECTURE

I. Introduction and Object of Formalization

The companion paper (Rhodes, 2026e; rhodesresearch.org) argues that the dominant failure mode in Western China analysis is input corruption rather than model deficiency, and proposes a five-channel triangulation framework augmented by three predictive components: a signal hierarchy, a silence map, and a contradiction index. That paper is discursive and applied; this one is formal. The object of formalization is fourfold: to derive the five-channel architecture from first principles via the Ledermann bound; to give the silence map a proper statistical treatment as informative censoring; to formalize prediction as a competing-risks hazard problem; and to present the first

empirical calibration of the framework using the March-June 2026 commodity shock as a natural identification instrument.

The formalization serves three purposes beyond rigor for its own sake. First, it converts informal rules — "three independent channels must converge," "silence is signal," "a widening gap predicts instability" — into statements with precise conditions, so that the framework's claims become checkable rather than rhetorical. Second, it exposes which of the framework's assumptions are load-bearing: the conditional independence of channel errors turns out to be the identifying assumption on which everything else rests, which is exactly the vulnerability the companion paper identifies informally as the coordination-attack scenario. Third, it makes the framework estimable. Every object defined below — latent states, loadings, lags, suppression parameters, hazards — has a posterior distribution given data, and the paper specifies the computational strategy for obtaining it.

II. Primitives and Notation

Let $(\Omega, \mathcal{F}, P)$ be a probability space. Time is discrete, $t = 1, \dots, T$. The latent state $x_t = (r_t, n_t)^T \in \mathbb{R}^2$ consists of behavioral reality r_t and the official narrative n_t , both unobservable. The contradiction index $C_t \triangleq n_t - r_t$ is the signed gap. Five channels $i = 1, \dots, 5$ correspond to: capital flows, business intelligence, diaspora and migration, official headline indicators, and the silence map. Each channel produces an observation y_{it} that may be continuous, ordinal, or censored. The complete data are (y_t, M_t) where $M_t = (M_{1t}, \dots, M_{5t})$ is the missingness vector, $M_{it} = 1$ indicating that channel i is suppressed at time t .

III. State Dynamics

III.1 Behavioral Reality

Behavioral reality follows a near-unit-root process:

$$r_t = \rho_r r_{t-1} + \eta_t, \quad \eta_t \sim (0, \sigma_r^2), \quad \rho_r \in (0.85, 1.00)$$

The near-unit-root specification captures the stylized fact that underlying economic and political conditions evolve slowly relative to narrative adjustments. The innovation variance σ_r^2 is identified by the assumption that at least one channel loads primarily on reality rather than narrative.

III.2 The Managed Narrative: An Optimizing Rule

The narrative follows a Guriev-Treisman optimizing rule:

$$\begin{aligned} n_t &= \omega_t n_{t-1} + (1-\omega_t) \theta_t + \varepsilon_t \\ \omega_t &= \Lambda(\alpha + \beta C_{t-1} - \gamma S_{t-1}) \end{aligned}$$

where θ_t is the period- t optimal narrative (the message that maximizes regime support given current conditions), S_{t-1} is the silence footprint (defined in Section VI), and $\Lambda(\cdot)$ is the logistic function. This rule endogenously generates hysteresis: when the contradiction index is wide (βC_{t-1} large), narrative inertia ω_t rises — the regime doubles down rather than corrects. When the silence footprint contracts (γS_{t-1} falls), inertia falls and correction becomes more likely. The narrative corrects late and abruptly, not smoothly, and suppression expands precisely when the gap is widest — the dynamic observed repeatedly in managed-information systems.

IV. The Measurement System

IV.1 Observation Equations

The latent signal for channel i at time t is:

$$y^*_{it} = a_i + b_i^r r_{t-l(i)} + b_i^n n_{t-l(i)} + u_{it}, \quad u_{it} \sim (0, \sigma_i^2)$$

where $l(i)$ is channel i 's characteristic lag and (b_i^r, b_i^n) are factor loadings. Capital flows and diaspora load primarily on reality ($b_i^n \approx 0$). Official headline indicators load primarily on narrative with high idiosyncratic variance ($b_i^r \approx 0$, σ_i^2 large). Business intelligence loads more heavily on reality. The silence map (channel 5) is the informative-censoring channel formalized in Section VI.

IV.2 The Signal Hierarchy as a Lag Prior

The signal hierarchy of the companion paper — capital flows lead by 6-18 months, business intelligence by 3-9 months, diaspora by 1-6 months, official indicators contemporaneous — is incorporated as a prior on the lag structure $l(i)$. This converts the informal hierarchy into an estimable parameter with uncertainty: the posterior distribution of $l(i)$ concentrates on the ranges suggested by the hierarchy but is not restricted to them. Evidence that a channel's lead time has shortened (an indicator of coordination attack) appears as posterior mass shifting toward shorter lags.

IV.3 Ordinal Measurement for Qualitative Channels

Several channels produce ordinal rather than cardinal observations: business intelligence reports, diaspora sentiment indices, silence footprint counts. For ordinal channel i with K ordered categories, the observation model is:

$$y_{it} = k \quad \text{iff} \quad \tau_{k-1} < y^*_{it} \leq \tau_k, \quad k = 1, \dots, K$$

where $\tau = (\tau_0, \dots, \tau_K)$ are threshold parameters with $\tau_0 = -\infty$ and $\tau_K = +\infty$. The thresholds are identified by normalization of the continuous latent signal.

IV.4 The Conditional-Independence Axiom

The identifying assumption of the entire framework is conditional independence of channel errors given the latent state:

$$u_{it} \perp u_{jt} \mid x_t \quad \text{for all } i \neq j$$

This axiom is load-bearing. It fails under coordination attack — the scenario in which multiple channels are simultaneously manipulated to produce consistent false signals. The companion paper's practical guidance for detecting coordination attack (look for implausible cross-channel agreement, absence of the natural variance that independent channels generate) translates, in the formal model, to posterior diagnostics: a coordinated-attack episode appears as anomalously low estimated channel error covariances, flagging an identification threat. The axiom cannot be tested from data alone; it requires substantive judgment about the independence of channel generating processes.

V. Identification: Why Five Channels, and Why Three-Channel Convergence

The Ledermann (1937) bound: for a k -factor model with p indicators, identification of both loadings and idiosyncratic variances requires:

$$p \geq 2k + 1$$

With $k = 2$ latent factors (behavioral reality and official narrative), $p = 5$ is the minimum. The five-channel architecture is not arbitrary — it is the minimal measurement system for the two-factor structure. Fewer channels leave the model under-identified even with strong priors on the loading signs.

The Anderson-Rubin (1956) three-indicator condition establishes that three channels are sufficient for single-factor identification, via the tetrad constraint:

$$\sigma_{12}\sigma_{34} = \sigma_{13}\sigma_{24}$$

This explains the informal "three-channel convergence" rule of the companion paper: convergence across three independent channels is sufficient evidence that a single underlying condition (behavioral reality) is being observed, because three channels over-identify the single factor. The formal derivation clarifies both why three channels are sufficient for detection and why five are required for the full two-factor structure.

VI. The Silence Map as Informative Censoring

VI.1 The Suppression Mechanism

Define censoring indicator $M_{it} = 1\{y_{it} \text{ is suppressed}\}$. Under ignorable missingness, the censoring mechanism is independent of the latent state: $P(M_{it} = 1 | x_t) = P(M_{it} = 1)$. This is the missing-completely-at-random assumption, and it is false in managed-information systems. The informative-censoring assumption:

$$P(M_{it} = 1 | x_t) = \Lambda(y_\theta + y_1 C_t), \quad y_1 > 0$$

Suppression probability is increasing in the contradiction index. When the gap between narrative and reality widens, the probability of suppressing channels that would reveal the gap rises. This is the mechanism the companion paper identifies informally: silence is not the absence of information but its most concentrated form.

VI.2 Why Silence Enters the Likelihood

Under informative censoring, the missingness mechanism must enter the likelihood. Rubin (1976) established that ignoring the missingness mechanism biases inference whenever missingness depends on the latent state. The complete-data likelihood factors as:

$$L(\theta | y_{obs}, M) = L(\theta | y_{obs}) \times P(M | x_t, y)$$

The second term is not ignorable. Estimators that condition only on observed channels — treating silence as the absence of data — are biased in the direction of overestimating behavioral reality, precisely because high-reality-gap periods are the periods from which data are most systematically missing. Correctly specifying the censoring mechanism turns this bias into information.

VI.3 The Silence Footprint

The silence footprint $S_t = \sum_i M_{it}$ is an ordinal instrument for C_t : under the informative-censoring assumption, $E[S_t | C_t]$ is increasing in C_t . This result — that a count of suppressed channels is an unbiased ordinal measure of the contradiction index — formalizes the companion paper's claim that "the pattern of what is not said is itself informative." The footprint enters both the likelihood

(through the censoring model) and the hazard (through its effect on narrative inertia in the optimizing rule).

VII. The Contradiction Index as a Filtered State

The estimated contradiction index is $\hat{C}_{t+} = \hat{n}_{t+} - \hat{f}_{t+}$, where hats denote Kalman-filtered means. In the full nonlinear model, filtering uses the particle Gibbs algorithm of Lindsten, Jordan, and Schön (2014), which handles the ordinal channels and informative-censoring silence map without linearization. The linear-Gaussian sub-block — appropriate when channels are cardinal and uncensored — uses standard Kalman recursions with Rauch-Tung-Striebel smoothing.

The filtered contradiction index has a natural interpretation: it is the posterior mean of the gap between what the official narrative claims and what behavioral reality supports, given all observed channel data and the correct treatment of suppressed channels as informative absences rather than missing observations.

VIII. Prediction: Competing Risks Over Three Terminal Regimes

Three terminal regimes exhaust the model's prediction space:

Regime 1 (Policy Reversal): narrative acknowledges the gap and policy adjusts toward behavioral reality. The necessary condition is a sustained decline in the contradiction index driven by reality improvement rather than narrative collapse.

Regime 2 (Escalating Control): suppression expands rather than narrative adjusts. The silence footprint rises while the contradiction index widens further. Historical examples include the 2023 suspension of youth unemployment reporting.

Regime 3 (Delayed Recognition): gap maintained until external shocks make suppression infeasible. The narrative is then forced to correct discontinuously — the Guriev-Treisman "spin dictator" endpoint in which the managed information system loses credibility faster than it loses control.

The cause-specific hazard for regime j :

$$\lambda_j(t \mid \hat{C}_{t+}, S_t) = \lambda_{j0}(t) \exp(\alpha_j \hat{C}_{t+} + \beta_j S_t)$$

Sign predictions: $\alpha_1 < 0$ (higher contradiction index reduces probability of voluntary reversal), $\alpha_2 > 0$ (higher contradiction index raises probability of escalating control), $\beta_2 > 0$ (higher silence footprint signals escalating control). These sign restrictions are testable and constitute the framework's falsifiable core.

IX. Application: The Yuan Internationalization Question as a Gated Event

IX.1 Structure

Genuine reserve currency status requires capital account convertibility as a necessary condition — modeled as an absorbing state $A = 1$. The two-stage structure:

$$P(\text{reserve status by } T) = P(A = 1 \text{ by } T) \times P(\text{credibility achieved} \mid A = 1)$$

A system with a widening C_t is moving away from the transparency conditions convertibility requires. A regime that cannot credibly commit to not suppressing inconvenient data cannot credibly commit to not suppressing inconvenient capital flows — the same information-architecture failure that generates the contradiction index also generates the credibility deficit that blocks convertibility.

IX.2 The Opening Hazard and the Silence Map

The hazard for capital account opening:

$$\lambda_A(t \mid \hat{C}_{t+}, S_t) = \lambda_{A_0}(t) \exp(-\delta \hat{C}_{t+} - \zeta S_t), \quad \delta, \zeta > 0$$

Both a wide contradiction index and a high silence footprint reduce the probability of opening, because both are structural features the model identifies as incompatible with the transparency conditions convertibility requires. The opening hazard is not identified from the Chinese case alone — a single censored spell — but is calibrated against the historical experience of partial liberalizations that reversed (Russia 1998, Argentina 2001, Turkey 2018).

IX.3 Channel Evidence in the Notation

The renminbi's persistent share of 2.7-3.1 percent of SWIFT-tracked global payments despite a decade of active promotion is the observable consequence of the gated structure: the credibility stage remains blocked. Each year of non-opening is an observation on the survival function, contributing to the likelihood of the opening hazard even as a censored spell. The model's 75-80 percent probability of non-occurrence over a five-year horizon falls out as the probability that the absorbing state constraint remains binding — consistent with the available data without requiring a structural break in Chinese capital account policy.

IX.4 Sensitivity of the Gated Prediction

The gated probability is sensitive to δ (the elasticity of opening hazard to the contradiction index) and robust to the credibility stage specification (whether credibility is a binary threshold or a continuous process). The sensitivity analysis identifies the conditions under which the 75-80 percent range shifts: a sustained decline in $\hat{C}_{t+}$ of more than 3 points over 24 months, driven by reality improvement rather than narrative collapse, would shift the five-year probability toward 50-60 percent. No such decline is visible in the current calibration window.

X. Falsification, Diagnostics, and Confidence Degradation

Four falsification conditions apply to the framework as a whole:

1. Three or more channels converge on $\hat{C}_{t+} \approx 0$ for a sustained period (eight or more consecutive quarters), indicating that narrative and behavioral reality have genuinely aligned.
2. Silence footprint contracts persistently while the contradiction index remains wide — the pattern inconsistent with the informative-censoring model, which predicts co-movement.
3. Terminal regime transition inconsistent with sign predictions: specifically, Regime 1 transition accompanied by rising rather than falling contradiction index, or Regime 2 transition accompanied by a falling silence footprint.
4. Genuine reserve currency status achieved without prior sustained capital account opening — which would require a credibility mechanism not captured in the framework.

X.1 Robustness to Channel Degradation

The framework degrades gracefully when channels become unavailable. With four channels (one suppressed), the Ledermann bound requires relaxing one identification constraint — typically imposing a prior that one loading is zero. With three channels, single-factor identification is maintained via the Anderson-Rubin condition, but the two-factor structure is under-identified; inference about the narrative factor specifically becomes unreliable. The companion paper's practical guidance — treat sustained silence as raising the contradiction index estimate, not as removing information — maps precisely onto the informative-censoring treatment: a degraded channel contributes to the likelihood through its suppression probability rather than through its observation.

Confidence degrades if channel covariance structure departs more than two posterior standard deviations from calibrated values for more than four consecutive periods — a signal that either the conditional-independence axiom is failing (coordination attack) or the factor structure has changed.

XI. The Strongest Counterargument

The most serious objection to this framework is not empirical but architectural: the conditional-independence axiom on which the entire identification strategy rests is unverifiable from data alone. If a sophisticated actor coordinates the manipulation of all five channels simultaneously — producing consistent false signals across capital flows, business intelligence, diaspora behavior, official indicators, and the silence map — the framework cannot distinguish that coordination from a genuine underlying condition. The Ledermann bound establishes that five channels are sufficient for identification given conditional independence; it provides no protection if that axiom fails.

This is a genuine vulnerability, and it deserves a direct answer rather than a hedge.

First, coordinated manipulation of all five channels is substantially more difficult than it appears. The channels are structurally heterogeneous: capital flows reflect the decisions of thousands of firms and individuals operating under legal systems outside Chinese jurisdiction; diaspora behavior reflects the private choices of people who have left; offshore credit spreads reflect assessments by counterparties with no obligation to the Chinese narrative. Coordinating these channels requires not just controlling the information environment but altering the behavior of actors who are, by definition, partially outside the coordination reach of the managed system. The companion paper identifies this as the framework's key practical protection: the channels were chosen precisely because their independence is structural, not merely assumed.

Second, the coordination-attack scenario is itself detectable, though not through the primary identification strategy. Anomalously low posterior estimates of channel error covariances — channels agreeing more than their structural independence would predict — flag a potential coordination event. The framework degrades to a warning rather than an estimate: it cannot tell you what the underlying condition is, but it can tell you that something unusual is suppressing the natural disagreement between channels. That warning is itself informative.

Third, the falsification conditions in Section X are specified as formal interventions on model parameters precisely so that a coordination-attack scenario that eventually unravels produces a clear disconfirmation signal. The framework earns its credibility not by being immune to

coordination attack but by specifying in advance what such an attack would look like in the data and what would constitute evidence that it had ended.

The honest position: this framework is more robust to partial manipulation than to full coordination attack, and more useful for detecting gradual deterioration than for identifying sudden reversals engineered across all channels simultaneously. Those are genuine limitations. They do not invalidate the framework's application to the current China case, where the structural heterogeneity of the channels and the observable behavioral evidence — firm exit decisions, migration patterns, offshore spread behavior — are not plausibly coordinated artifacts.

XII. Conclusion: The Formal Architecture

The five-channel framework is the minimal measurement system for identifying a two-factor reality-narrative structure, derived from the Ledermann bound. The silence map is informative censoring — not missing data but a censoring mechanism that enters the likelihood and carries information about the contradiction index. The contradiction index is the filtered latent gap. Prediction is a competing-risks hazard problem with testable sign restrictions and explicit falsification conditions.

What the formalization establishes beyond the companion paper's informal claims: the framework is estimable, the identifying assumptions are explicit, the degradation conditions are specified, and the falsification conditions are formal interventions on model parameters rather than verbal statements. The second part of this paper implements the linear-Gaussian sub-block of the model and demonstrates what the framework delivers empirically when a natural identification instrument arrives.

* * *

PART II: EMPIRICAL CALIBRATION

XIII. The Identification Problem and the Commodity Shock Instrument

The theoretical model faces a chronic identification obstacle: with both loadings and idiosyncratic variances unknown, the latent scale is set only by normalization, and the two factors are separable only where some observable moves one and not the other. In ordinary data, the two factors co-move — the narrative tracks reality with a managed gap — so the separation is unavailable.

The March-June 2026 commodity shock provides the required separation. The Iran conflict raised world commodity and energy prices through a channel independent of Chinese domestic demand. Three identifying properties establish the shock as a valid instrument: it is exogenous (world-priced, orthogonal to Chinese domestic conditions); it is observed (mining-stage price series corroborated by Tier 1 commodity data, entering the model as a measured regressor, not a latent); and its arithmetic is known (mechanical propagation through the producer-price system is calculable stage by stage). Everything in the price data that the shock explains is neither reality nor narrative. What loads on r_t is the price system's residual behavior at the household-facing stage; what loads on n_t is the condition the official release structure narrates. The shock is, in the estimation sense, an instrument — and it arrived without being asked for.

A design decision deserves its paragraph. The headline PPI channel is not purged of the shock's mechanical component before loading on n_t . The narrative factor is defined as the narrated condition — the state of the economy a reader of the official release structure would infer — and the regime narrates the headline. That the headline's improvement is mechanically borrowed from world commodity arithmetic is not a measurement contaminant to be removed; it is the phenomenon itself, the mechanism by which the gap opens. The shock's influence on n_t is routed through the narrative transition ($\gamma_n \xi_t$), where it belongs conceptually: the shock changed what could be narrated. This is the Guriev-Treisman harvesting mechanism made operational.

XIV. The Linear-Gaussian Sub-Block

The full companion model is nonlinear and non-Gaussian, requiring particle Gibbs estimation. For the January-June 2026 calibration window, the active channels are approximately continuous and the commodity instrument is strong enough to identify the factor separation without the full particle filter. The linear-Gaussian sub-block suffices for anchoring the scale and location of the factors.

The state vector is $f_t = (r_t, n_t, r_{t-1})'$, the lag carried to serve the unemployment channel's one-month delay. Transitions:

$$\begin{aligned} r_t &= \varphi_r r_{t-1} + \eta_{r,t} \\ n_t &= \psi_n n_{t-1} + \gamma_n \xi_t + \eta_{n,t} \end{aligned}$$

The exogenous shock enters the narrative transition, not the reality transition. The reality transition receives no shock term because the maintained hypothesis under test is precisely whether the shock reaches the real demand condition; putting it there would assume the answer. Measurement equations, with units and normalizations:

$$\begin{aligned} c_t &= \delta c \xi_t + r_t + \varepsilon c && (\text{consumer-goods PPI; } \lambda = 1 \text{ sets the scale of } r) \\ u_t &= \mu u - \lambda u r_{t-1} + \varepsilon u && (\text{youth unemployment, one-month lag, inverse}) \\ m_t &= \mu m + \lambda m r_t + \varepsilon m && (\text{PMI output-minus-input price spread}) \\ h_t &= n_t + \varepsilon h && (\text{headline PPI as narrated; } \lambda = 1 \text{ sets the scale of } n) \\ p_t &= \mu p + \lambda p n_t + \varepsilon p && (\text{industrial profit headline}) \end{aligned}$$

Measurement noise is tier-weighted per the admissibility hierarchy: the consumer-goods channel carries the smallest variance; the profit headline, a Tier 4 level series with known administered-price content, carries the largest; unemployment sits between, treated as a managed lower bound whose changes are more informative than its level.

XV. Data Panel, Calibration, and Results

The single externally calibrated structural parameter is the baseline pass-through δc — the fraction of a mining-stage price change reaching consumer-goods producer prices within the year under pre-existing demand suppression. The anchor is the 2021 episode: World Bank contemporaneous analysis documents downstream PPI holding roughly flat near +0.5 percent year-on-year through

the 2021H2 commodity boom while upstream stages surged. That anchor places δc in the low single digits; the central calibration uses 0.03, with the sensitivity grid spanning 0.00 to 0.06.

Remaining parameters (persistence $\phi_r = 0.95$, narrative inertia $\psi_n = 0.90$, harvesting coefficient $\gamma_n = 0.10$, channel loadings and noise) are calibrated, not estimated, and tabulated in Appendix B. Six months of data cannot support maximum-likelihood estimation of a two-factor system. What six months containing an exogenous shock can support is exactly what is claimed here — anchoring the location and scale of the factors and the index built on them.

XIV.1 Smoothed Factor Paths

Central calibration results ($\delta c = 0.03$):

January: $r_t = -1.29$ (sd 0.18), $n_t = -0.63$ (sd 0.40), $C_t = 0.65$ (sd 0.44)

February: $r_t = -1.32$ (sd 0.14), $n_t = -0.22$ (sd 0.34), $C_t = 1.10$ (sd 0.37)

March: $r_t = -1.26$ (sd 0.14), $n_t = +0.64$ (sd 0.33), $C_t = 1.90$ (sd 0.36)

April: $r_t = -1.21$ (sd 0.14), $n_t = +2.20$ (sd 0.33), $C_t = 3.41$ (sd 0.36)

May: $r_t = -1.25$ (sd 0.15), $n_t = +3.55$ (sd 0.34), $C_t = 4.80$ (sd 0.37)

June: $r_t = -1.28$ (sd 0.17), $n_t = +4.62$ (sd 0.39), $C_t = 5.90$ (sd 0.43)

Units: producer-price year-on-year points of each factor's defining channel. Smoothed via Rauch-Tung-Striebel recursion.

XIV.2 Three Results in Ascending Order of Consequence

First, the behavioral-reality factor is statistically motionless: r_t ranges from -1.32 to -1.21 across the six months — a band narrower than its own posterior standard deviation — through the largest exogenous input-cost shock since 2021. The latent demand condition of the household-facing economy in June, mid-shock, is indistinguishable from its condition in January, pre-shock.

Second, the narrative factor absorbed the shock nearly in full: n_t rises 5.2 points from January to June, tracking the headline the release structure narrates. The decomposition of that rise is the point: it is borrowed, mechanically, from year-on-year commodity arithmetic, and the regime's attribution of it to the anti-involution campaign converts an exogenous windfall into a claimed policy vindication.

Third, the contradiction index expands from 0.65 to 5.90 — nine-fold in six months, a widening of more than eleven posterior standard deviations. A C_t of roughly 0.65 in January describes narrated and lived price conditions in ordinary disagreement; a C_t approaching 5.90 in June describes an official economy reflatting at 4 percent while the household-facing economy deflates at one — conditions that do not merely differ in degree but in sign.

XIV.3 Sensitivity Grid

Across $\delta c \in \{0.00, 0.03, 0.06\}$, June C_t spans only 5.62 to 6.19, because the shock-adjustment term $\delta c \xi_t$ moves r_t by at most ± 0.28 points at the grid extremes. The finding is not an artifact of the pass-through calibration. The minimum contradiction index across all specifications is 3.71 — well above the threshold at which the competing-risks hazard layer activates.

XVI. The Deterministic Decay and the Hazard Calendar

What elevates this calibration above a descriptive exercise is that the narrative's support decays on arithmetic. Year-on-year statistics have a memory of exactly twelve months: the mining-stage price level that surged in April-June 2026 becomes the comparison base in April-June 2027, and unless world commodity prices rise again — a second exogenous event, not a policy option — the narratable headline mean-reverts on a schedule fixed by the calendar.

Three commodity-level scenarios:

Flat at June 2026 levels: implied headline PPI crosses zero in May 2027. Central case — the borrowed component expires with the base effect.

−10% retracement (post-ceasefire normalization): headline crosses zero in March 2027. The June 2026 monthly print (−0.3% MoM) is the leading edge of this path.

+10% further appreciation: headline crosses zero beyond June 2027. Requires a new exogenous shock; not a policy instrument Beijing controls.

The forced narrative-correction window is therefore March-June 2027. Inside that window, the number the narrative currently leads with reverts toward or below zero while — on the reality factor's demonstrated insensitivity — the household-facing condition it papered over remains where it was. The Guriev-Treisman closure then applies with unusual precision: the narrative rule maintains gaps until an exposure threshold is crossed, and here the exposure event is not a stochastic leak but the arithmetic of the regime's own chosen indicator.

The reality factor did not move. The narrative factor rose five points on borrowed arithmetic. The loan comes due between March and June 2027.

The gap must be resolved within the window. The competing-risks hazard layer enumerates the only three resolutions: policy reversal (genuine demand-side stimulus repricing the household economy — rated structurally obstructed by the political-economy analysis in the companion papers), escalating control (methodology revision, indicator substitution, or suppression of the newly inconvenient series — the 2023 youth-unemployment playbook, whose marginal cost rises with each use), and delayed recognition (narrating the reversion as "high-base effects," technically true and cumulatively corrosive). This calibration does not estimate the hazard split — that requires the full model — but delivers what the hazard layer lacked: a gate date. The correction event is near-certain on arithmetic; only its form is in question.

XVII. Falsification Conditions for the Dated Prediction

A framework that writes down a March-June 2027 window for a forced narrative event must specify in advance what would disconfirm it. Three terminal regimes, three falsification sets:

Regime 1 (Policy Reversal) confirmed if: sustained decline in C_t driven by r_t improvement (not n_t collapse); capital account liberalization steps announced and implemented; private investment returns to positive year-on-year growth for three or more consecutive quarters.

Regime 2 (Escalating Control) confirmed if: headline PPI methodology revision that severs the base-effect arithmetic; suspension or redefinition of one or more channels used in this calibration; silence footprint rises while officially reported conditions improve.

Regime 3 (Delayed Recognition) confirmed if: headline PPI reverts to negative without policy acknowledgment; official communication attributes reversion to "base effects" or "external normalization" without demand-side policy response; contradiction index remains wide through the correction window.

Disconfirmation: the dated prediction is falsified if the headline PPI does not cross zero in the March-June 2027 window under any of the three regime characterizations — specifically, if commodity prices appreciate sufficiently to extend the narratable positive headline beyond June 2027 without a new exogenous shock. This would require world commodity prices to remain elevated or rise further, which the model treats as a low-probability scenario given the Iran conflict's documented trajectory, but which it cannot rule out. A second commodity shock that extends the window is not a falsification of the framework; it is a new identification event.

XVIII. Generalizability: Historical Back-Tests

The framework is applied qualitatively to three historical cases to establish generalizability beyond the China application. These are directional back-tests, not formal re-estimations.

Russia 2014-2022: Capital flow channel (ruble positions, reserve drawdown trajectories) produced consistent behavioral-reality signals from 2014 onward that diverged from official narrative of resilience and recovery. The contradiction index would have widened continuously from 2015 through 2022. Terminal regime: Delayed Recognition, forced by the February 2022 shock. Framework performance: strong on identification, weak on terminal-regime discrimination — the competing-risks split between Escalating Control and Delayed Recognition was not resolvable from available channels prior to February 2022.

Iran 2019-2021: Framework identifies widening contradiction index from capital flight and unofficial exchange rate divergence. Terminal regime: Delayed Recognition via external pressure rather than internal correction. Framework limitation: channel degradation under sanctions reduces the cardinal channels available, leaving heavy reliance on the silence map and diaspora channels — both ordinal. The Ledermann bound is met, but with lower precision than in the China calibration.

Argentina 2018-2019: Capital flow channel provided clean, timely, directionally consistent behavioral-reality signals. The terminal regime (Delayed Recognition via external shock) was identifiable in advance from the hazard structure. The framework's limitation in this case is the soft silence footprint — the informative-censoring formalization works most cleanly with hard suppression rather than the softer selective-emphasis patterns that characterize democratic or semi-democratic managed narratives.

Synthesis: across the three back-tests, the framework demonstrates consistent identification of widening contradiction indices from behavioral channel data in advance of official acknowledgment. Its strongest performance is in cases with rich capital flow data and hard suppression signals. Its genuine limitations are terminal-regime discrimination and channel degradation under sanctions.

XIX. Limitations and the Path to Full Estimation

Three limitations bound the empirical results of Sections XII-XV.

First, the calibration implements the linear-Gaussian sub-block rather than the full nonlinear non-Gaussian model. The ordinal channels and informative-censoring silence map require the particle Gibbs algorithm of Lindsten, Jordan, and Schön (2014). The linear-Gaussian results are valid within the January-June 2026 window; outside it, the full model is required. The quantitative difference between the sub-block result (nine-fold widening) and the full model result (which incorporates all ordinal channels) would adjust the cardinal estimates but is not expected to reverse the directional finding given the sensitivity grid results.

Second, the pass-through parameter $\delta c = 0.03$ is calibrated against a single historical episode (2021). The sensitivity grid demonstrates robustness across the calibrated range, but a wider range of historical episodes would sharpen the estimate.

Third, the back-tests in Section XVII are qualitative applications rather than formal re-estimations. Formal back-testing requires constructing the data panels for each historical case and re-running the filter. The qualitative applications establish directional plausibility; the formal back-tests remain for future work.

The March-June 2027 window serves double duty: it is the model's first dated out-of-sample prediction and, when it resolves, a third identifying event for the estimation that follows. The model predicts its own next calibration point.

XX. Conclusion

This paper makes four contributions. First, it provides the formal mathematical architecture for the five-channel framework: the Ledermann identification argument, the informative-censoring silence map, the filtered contradiction index, and the competing-risks hazard prediction system with testable sign restrictions. Second, it implements the linear-Gaussian sub-block on January-June 2026 data, exploiting the commodity shock as a clean identification instrument, and reports a nine-fold widening of the contradiction index robust across the full sensitivity grid. Third, it derives a calendared prediction — a forced narrative-correction window in March-June 2027 — from the deterministic decay of the narrative's borrowed support, and maps falsification conditions for that prediction to the three terminal regimes. Fourth, it applies the framework to three historical back-test cases, with honest assessment of where the framework succeeds and where it faces genuine identification challenges.

The prediction is the most testable claim in the portfolio. It specifies a window, a mechanism, three possible outcomes, and the observable evidence that would distinguish them. The framework earns its credibility not by getting the prediction right — that is a function of the underlying conditions — but by stating in advance what right and wrong look like, and honoring that commitment when the window arrives.

* * *

Appendix A: Proof Sketches

A.1 Ledermann Bound ($k = 2, p = 5$)

The number of free parameters in (β, Σ) is $p(k+1) - k(k-1)/2$. Identification requires $p(p+1)/2 \geq p(k+1) - k(k-1)/2$, simplifying to $p \geq 2k+1$. For $k = 2$: $p \geq 5$. $\square$

A.2 Three-Indicator Solution (Anderson and Rubin, 1956)

With three channels and one factor, the pairwise covariances $\sigma_{ij} = b_i b_j \sigma^2$ identify loadings via $b_1^2 = \sigma_{12}\sigma_{13}/\sigma_{23}$ and cyclic permutations. With two channels, one covariance constrains two unknowns: under-identified. $\square$

Appendix B: Calibrated Parameters

φ_r (reality persistence): 0.95 — structural condition; near-unit-root within window

ψ_n (narrative inertia): 0.90 — release-structure stickiness

γ_n (narrative harvesting of ξ): 0.10 — matches headline swing given ψ_n ; over-identified once the 2021 event is added

δc (baseline pass-through): 0.03 [0.00, 0.06] — World Bank 2021H2 anchor

λ_u (unemployment loading, lag 1): 1.2 (inverse) — scale-matched to 2025H2 comovement

λ_m (PMI spread loading): 2.5 — survey-diffusion to pp conversion

λ_p (profit-headline loading): 4.0 — headline profit variance ratio; largest noise assigned

Measurement sd (c, u, m, h, p): 0.30, 0.45, 1.50, 0.60, 5.00 — tier-weighted per admissibility hierarchy

Innovation sd (r, n): 0.15, 0.40 — reality structural; narrative discretionary

Initial state (r_0 , n_0 , P_0): -1.5, -1.2, diffuse — late-2025 deflation levels

The filter is a standard Kalman recursion with Rauch-Tung-Striebel smoothing and row-wise deletion for missing observations; the state is augmented with r_{t-1} to carry the unemployment channel's one-month lag. The complete reproduction script (Python, NumPy only for the filter, Matplotlib for figures) is available as `calibrate_latent.py`; every result regenerates from it deterministically.

References

- Anderson, T. W., & Rubin, H. (1956). Statistical inference in factor analysis. *Proceedings of the Third Berkeley Symposium on Mathematical Statistics and Probability*, 5, 111-150.
- Guriev, S., & Treisman, D. (2022). *Spin Dictators: The Changing Face of Tyranny in the Twenty-First Century*. Princeton University Press.
- Ledermann, W. (1937). On the rank of the reduced correlational matrix in multiple-factor analysis. *Psychometrika*, 2(2), 85-93.
- Lindsten, F., Jordan, M. I., & Schön, T. B. (2014). Particle Gibbs with ancestor sampling. *Journal of Machine Learning Research*, 15, 2145-2184.
- Reinhart, C. M., & Rogoff, K. S. (2009). *This Time Is Different: Eight Centuries of Financial Folly*. Princeton University Press.

Rhodes, J. K. (2026a). The Good, the Bad and the Ugly: Dollar Dominance, Institutional Failure, and the Architecture of De-Dollarization. Rhodes Research Working Paper. rhodesresearch.org.

Rhodes, J. K. (2026b). Reality Leaves Footprints: How to Understand China When Every Traditional Source Is Compromised. Rhodes Research Working Paper. rhodesresearch.org.

Rhodes, J. K. (2026c). The Company Store: Why the RMB Is Built for Lock-In, Not Reserve Status. Rhodes Research Working Paper. rhodesresearch.org.

Rhodes, J. K. (2026d). The Efficiency Trap: How the World Drifted Into Chinese Dependency. Rhodes Research Working Paper. rhodesresearch.org.

Rhodes, J. K. (2026e). From Triangulation to Prediction: A Five-Channel Framework for Forecasting Chinese Policy Outcomes. Rhodes Research Working Paper. rhodesresearch.org.

Rhodes, J. K. (2026f). The Multiplier Effect and Its Mirror Image: Why Each Foreign Firm That Leaves China Triggers a Contractionary Chain Reaction Larger Than the Departure Itself. Rhodes Research Working Paper. rhodesresearch.org.

Rubin, D. B. (1976). Inference and missing data. *Biometrika*, 63(3), 581-592.

World Bank. (2022). *Commodity Markets Outlook: The Impact of the War in Ukraine on Commodity Markets*. Washington, DC: World Bank Group.

* * *

About the Author

J. Kenneth Rhodes is an independent researcher and writer. Former Chief of Securities, State of Tennessee; member of the Regulation D drafting committee; federal court expert witness in fraud cases; M.A. Economics. His research is available at rhodesresearch.org. Contact: jkr@rhodesresearch.org.