CPAF v6 — Technical Specification

A reproducible specification of the model's variables, equations, parameters, and update logic.

Based on :

Bilel ABBOUDI · Master of Arts in Public Policy (MAPP, 2024–25) · Jindal School of Government and Public Policy, O.P. Jindal Global University

O.P. Jindal Global University · 2026

Contents

A reproducible specification of the model's variables, equations, parameters, and update logic, extracted directly from the source (`config.js`, `simulation.js`) and verified at runtime. Written so a modeller could reconstruct or audit CPAF. Every equation is tagged with its grounding and linked to the dimension and theory it implements.

1. Purpose, scope, and how to read this

CPAF is a discrete-time system-dynamics model with a monthly step, coupling seven analytical dimensions (D1–D7) at runtime. It is a structured-reasoning instrument, not an empirically-calibrated forecaster: it reports reliability and ranges, not point predictions. This document specifies *what the model computes*; the Theoretical Foundations Companion specifies *which theory each mechanism embodies*, and the two are cross-referenced.

Provenance tags. Each variable, equation, and parameter carries one tag, matching the model's own four-tier scheme:

- **[E] Empirical** — from a documented measurement or dataset.
- **[M] Mechanism** — follows from an accounting identity or definitional relationship.
- **[L] Literature** — the relationship and its direction are supported by published research; the exact value is a reasoned quantification.
- **[A] Assumption** — a declared, adjustable value or equation form chosen for plausibility and theoretical fit.

Across all parameters the split is approximately **3% [E], 9% [M], 27% [L], 61% [A]**. Every parameter is visible and adjustable in the Admin panel; outputs are only as strong as these choices.

2. Architecture and the two coordination streams

One monthly `step()` advances the clock, then updates the stocks through their flow equations, classifies the regime, and updates coordination, politics, segments, and shocks (full order in §10). Seven dimensions are coupled at runtime, so each feeds the others within a step.

A keystone of the architecture is that **coordination is two distinct variables**, not one:

- **Artisan co-op** `coordination ∈ [0,1]` — a producer-side lever, built by cluster investment past a critical mass (§7.1). **[L]** Ostrom (1990); Granovetter (1978).
- **Inter-ministry delivery coherence** `policyCoherence` (C_eff) — a state-side multiplier applied to *every* policy effect (§6). **[L]** Peters (2015); Bardhan (2002).

The delivery multiplier used in the pipeline is `_coordEffect = (1 − coordRange) + coordRange × policyCoherence`, with `coordRange = 0.70` **[A]** — i.e. `_coordEffect = 0.30 + 0.70 × policyCoherence`.

3. State variables (stocks)

Each stock, with unit, range, initial values for the two cases, and grounding. Initial values are the `CASE_STUDIES[case].initial` block.

Stock	Symbol	Unit	Chikankari init	Cheriyal init	Grounding
Artisan income	`income`	₹/month	4,100	5,000	**[E]** source-informed (demonstration)
Artisan population	`population`	persons	250,000	20	**[E]** source-informed (demonstration)
Skilled artisans	`skills`	index 0–1	0.55	0.75	**[A]** illustrative
Market reputation	`market`	index 0–1	0.30	0.25	**[A]** illustrative
Middlemen control	`middlemen`	index 0–1	0.80	0.05	**[L]** Wilkinson-Weber (1999); Liebl & Roy (2003)
Artisan co-op	`coordination`	index 0–1	0.30	0.45	**[A]** illustrative; **[L]** Ostrom (1990)
Heritage capital	`heritage`	index 0–1	0.70	0.85	**[L]** UNESCO (2003) framing
Political feasibility	`political`	index 0–1	0.50	0.55	**[A]** illustrative (NPF)
Delivery coherence	`policyCoherence`	index 0–1	derived	derived	**[L]** Peters (2015)
Production capacity	`capacity`	units/mo	derived	derived	**[M]** = population × skills × infrastructure

Case parameterisation is for demonstration only. The two cases are illustrative instances that exercise the framework's mechanisms; their starting values are source-informed, not validated empirical calibrations or forecasts. The Chikankari starting income of ₹4,100/month is the average monthly income from Chikankari reported by Singh and Shah (2023); population ≈250,000 (Singh, 2020). Cheriyal income ₹5,000 and population ≈20 are from field and news sources (2020–2025).

Singh, R., & Shah, P. (2023). Work profile and occupational hazards of Chikankari workers with special reference to Lucknow, North-Eastern India. *The Pharma Innovation Journal, 12*(1), 1276–1279.

4. Thresholds and regime classification

Income lines (`THRESHOLDS`). TRAP = `2000` ₹/mo (**[E/L]** World Bank $3.00/day 2021 PPP ≈ Rangarajan rural line); GROWTH = `4122` ₹/mo (**[E]** rural MPCE, HCES 2023-24). The destitution floor `subsistence = 1450` (**[E/L]** World Bank $2.15/day 2017 PPP) is the income floor (§5).

Regime states. TRAP → TRANSITION → GROWTH, classified each step by `checkAttractor()`. Confirmation is **duration-gated (hysteresis)** — a single crossing is not a transition. **[L]** Scheffer (2009).

- Durations: TRAP_DUR = 6, GROWTH_DUR = 12, FALL_DUR = 3 months **[A]**; escape is asymmetric via `escapeMultiplier = 1.5`, so effective durations are `ceil(DUR × 1.5)` **[A]**.
- TRAP → TRANSITION when income has been above TRAP for ≥ `ceil(6 × 1.5)` months.

- TRANSITION → GROWTH when income has been above GROWTH for ≥ ceil(12 × 1.5) months.
- TRANSITION → TRAP when income has been below TRAP for ≥ 3 months.
- GROWTH → TRANSITION when income has been below 0.8 × GROWTH for ≥ 3 months.
- Near the trap boundary a **pull force** (pullStrength = 0.025 **[A/L]**, Scheffer 2009) resists escape.

5. The income model

Income is a **level set by the current stocks each step, not a compounding accumulator** — this removes both runaway growth and an artificial ceiling, and enforces the poverty trap as an AND-gate. The operative equation (stepIncome, verified) is:

```
income = subsistence + headroom × quality × share
```

where quality = clamp((skills + market) / 2, 0, 1) and share = clamp(1 − middlemen, 0, 1), then multiplied by monthly noise (0.98 + 0.04 × rng()) and floored at subsistence. **[A]** form is a declared level model; **[L]** the AND-gate (quality *and* a broken middleman) reflects monopsony extraction (Wilkinson-Weber 1999).

- subsistence = 1450 **[E/L]**, headroom = 8000 **[A]** (the earnable income above the floor).
- Because income depends on the skills, market, and middlemen stocks, **policy affects income indirectly** — by moving those stocks (§7), not by adding to income directly.
- Conditional reinforcing-loop activations (R1 skill→income, R2 market→income) gate the premia by state — e.g. below a subsistence band the skill premium is suppressed, and high middlemen capture it — implementing the trap's self-reinforcement. **[L]** Azariadis & Stachurski (2005).

6. The policy-effect pipeline

No policy acts at full strength. For each active, non-umbrella policy, and each affected stock k, calculatePolicyEffects() computes (verified):

```
effect[k] += policy.effects[k] × coverage × C_eff × ramp(t) × narrative × politicalCycle × prereq
```

then pairwise policyInteractions multipliers are applied. Each factor:

- **coverage** — the share of population the policy reaches **[A]**.
- **C_eff** = _coordEffect = 0.30 + 0.70 × policyCoherence — delivery coherence **[L]** Peters (2015).
- **ramp(t)** — sigmoid ramp-up to maturity, then **fatigue**: rampMonths = 36, initialEffectiveness = 0.10, fatigueStartYear = 4, fatigueRate = 0.06/yr **[L]** Baumgartner & Jones (1993); Kingdon (1984).
- **narrative** — ±alignment with the prevailing policy narrative **[L]** Shanahan et al. (2018).
- **politicalCycle** — electoral attention cycle: cycleLengthMonths = 60, attentionAmplitude = 0.18, implementationDip = 0.25 **[L]** Kingdon (1984).

- **prereq** — sequencing: a matured prerequisite boosts the effect (diminishing returns in the count met); none → penalty. **[L]** Howlett & Rayner (2007); Pierson (2000).

C_eff itself: `policyCoherence = clamp(0.4 + Σ coherenceContribution(active) − 0.02 × active_count, 0.1, 1)`, where per-policy `coherenceContribution` comes from a lookup (SFURTI 0.15, CLUSTER_DEV 0.15, TIES 0.10, NHDP 0.08, others 0) **[A]**. Coherence-building instruments offset the per-policy fragmentation penalty, so delivery need not fall as a portfolio grows.

7. Stock update equations

7.1 Artisan co-op (`updateCoordination`, verified). Built via a **critical-mass sigmoid**, not linearly:

`coordInvestment = Σ (policy.effects.coordination × coverage)` over non-umbrella policies;

`target = 1 / (1 + exp(−steep × (coordInvestment − crit)));`

`coordination += (target − coordination) × build`, clamped `[0.05, 0.95]`.

Parameters `crit = 0.05`, `steep = 60`, `build = 0.02` **[A]**. An **Olson inequality drag** lowers the reachable target when the top/bottom segment income ratio exceeds a threshold. **[L]** Granovetter (1978); Olson (1965).

7.2 Other stocks (governing forms; full flow terms in `stepSkills/Market/Population/Heritage/Middlemen`):

- `dSkills = training − erosion − exit`; `training_rate = 0.02`, `erosion_rate = 0.01` **[A]**. **[L]** human-capital.
- `dMarket = marketing + quality − erosion`; `growth_rate = 0.005`, `quality_effect = 0.10` **[A]**.
- `dPopulation = entry − exit`, income-differential driven, regime-dependent rates (TRAP entry 0.01/exit 0.05; GROWTH entry 0.05/exit 0.01) **[A/L]** Harris & Todaro (1970).
- `dHeritage = transmission − erosion − mass-production`; `transmission_rate = 0.02` **[A]**. **[L]** UNESCO (2003).
- `dMiddlemen = f(coordination, market, income, policies)` — broken by the co-op × market-channel AND-gate (`break_strength = 0.75`, `break_rate = 0.05`), with adaptive resurgence. **[L]** Galbraith (1952) countervailing power; Wilkinson-Weber (1999).

7.3 Political feasibility & narrative (`updatePolitical`, annual at month 12). `political` shifts with policy outcomes (welfare/market × improved), inequality, coordination, and regime, is amplified near elections, mean-reverts to 0.5, is dampened at the extremes, and clamps `[0.05, 0.95]`. **[L]** Kingdon (1984); Sabatier (1988); NPF.

8. The complexity layer (engineered stochasticity)

Devices that make the model non-stationary and complex (all admin-toggleable; `COMPLEXITY.enabled`):

- **Stochastic shocks** (`processComplexityShocks`, verified). Monthly probability `freq = baseFrequency (0.04)`, scaled by regime — `× trapVulnerability (1.8)` in TRAP, ×

`growthVulnerability (0.5)` in GROWTH — and raised by network density (`fragilityCoupling`, **[L]** Perrow 1984). On trigger: `magnitude = 0.05 + rng() × 0.15`; `income × (1 − magnitude)`, floored; secondary hits to `market` or `coordination` by type (monsoon, demand slump, input-cost spike, policy discontinuity, supply chain); recovery over `ceil(6 + magnitude × 12)` months at `recoveryRate = 0.4`. **[A]** values; **[L]** non-stationarity.

- **Demand saturation** — quadratic resistance above `saturationIncome = 10000` (above the model's typical ceiling, so rarely binding) **[A/L]** Kotler & Keller (2016).
- **Cross-dimensional tensions (T1–T7)** — heritage↔market, middleman adaptation (`adaptationRate 0.40`), coordination friction (`frictionPerPolicy 0.025`, `maxFriction 0.12`), brain drain (`drainRate 0.006`), political cycles, population competition (`baseIntensity 0.065`), inequality→coordination. **[L]** the corresponding theories in the companion.

9. Parameters (EQUATION_PARAMS) with provenance

Governing constants, each with value, admissible range, and tier. All are **[A]** unless noted; the model labels most "illustrative, calibratable."

Group	Parameter	Value	Range	Tier
income	subsistence (floor)	1450	1000–4000	**[E/L]**
income	headroom	8000	2000–20000	**[A]**
income	skill_premium	0.15	0.05–0.30	**[A]**
income	market_premium	0.12	0.05–0.25	**[A]**
income	coord_multiplier (coordRange)	0.70	0.50–0.90	**[A]**
income	middlemen_extraction	0.65	0.30–0.85	**[L]**
skills	training_rate	0.02	0.01–0.05	**[A]**
skills	erosion_rate	0.01	0.005–0.03	**[A]**
market	growth_rate	0.005	0.001–0.02	**[A]**
coordination	crit_mass	0.05	0.02–0.12	**[A]**
coordination	threshold_steep	60	20–120	**[A]**
coordination	build_rate	0.02	0.005–0.05	**[A]**
coordination	break_strength	0.75	0.30–0.90	**[L]**
thresholds	TRAP	2000	—	**[E/L]**
thresholds	GROWTH	4122	—	**[E]**

10. The update cycle (one `step()`)

Order of operations per month (verified from `step()`):

1. Advance the clock (`month`, `year`, `totalMonths`).

2. Compute `policyCoherence` (C_eff) and `_coordEffect` — the two-stream keystone.

3. `applyDelayedEffects()` — release queued (lagged) policy effects.

4. `processShocks()` then `processComplexityShocks()` — external and stochastic shocks.

5. `updateInfrastructure()`.

6. Stocks in order: `stepIncome → stepSkills → stepMarket → stepPopulation → stepHeritage → stepMiddlemen`.

7. `checkAttractor()` — regime classification and hysteresis.

8. `updateCoordination()` — the co-op sigmoid.

9. At month 12 only: `updatePolitical()` (and NPF) — annual political/narrative update.

10. `updateSegments()`, `checkTipping()`, then `computeNetworkDensity()` for the next step.

11. Stochasticity and reproducibility

Randomness is a seeded **mulberry32** generator (`rng()`), initialised to seed `12345` and **reset at the start of every run** (`initState`). The model is therefore stochastic *and* reproducible: a given seed and configuration reproduce a run exactly, which is what makes A/B comparison and calibration valid. Monte Carlo replicates a configuration under this noise to report reliability and range (`runMonteCarlo`).

12. Validation

CPAF is validated as a theory-driven structural model, using the subset of the Barlas (1996), Forrester–Senge (1980) and Sargent (2013) tests appropriate to a model with declared parameters; tests requiring empirical time-series or predictive fit are deliberately excluded, as they do not apply to a reasoning instrument. The suite runs in three stages across all three regimes, so a pass is a property of the structure rather than one calibration. **Result: 15/15 structural tests and 6/6 demonstration tests — 21/21.**

- **Stage 1 — direct structure (R1–R6):** mass conservation, variable boundedness, and the correct signs of extraction, complementarity, synergy, and fragmentation friction.
- **Stage 2 — structure-oriented behaviour (B1–B7, C1–C2):** multistability, threshold classification, ordered transitions, duration-gated hysteresis, gradual (non-instant) response, bounded convergence, ±15% stochastic robustness (B7), the coordination critical-mass threshold, and robustness across the admin parameter grid.
- **Stage 3 — behaviour pattern (S1–S5, C3):** poverty-trap persistence, integrated portfolio beating single-stance policy, the connective-failure case responding to market linkage, the distributional-failure case responding to structural policy, archetype divergence by lever, and the load-bearing-core portfolio structure.

The test-by-test record is in the Validation Report.

13. Sensitivity and Robustness

Because approximately 61% of parameters are declared assumptions, the model's *conclusions* are stress-tested so they do not depend on any single soft number. A one-at-a-time (OAT) sweep takes

each of the 28 ranged parameters to both extremes of its own declared range (N = 3 ensemble, 30-year horizon), and five core conclusions are re-checked at each setting. **Each conclusion holds across 55 of 56 parameter-extremes — 98% robust.** The few breaks occur only at degenerate values (e.g. a destitution floor set above the poverty line, or an earning ceiling below subsistence) that no reasonable calibration would adopt. The global complement is validation test B7 (a ±15% all-parameter Monte Carlo).

The five conclusions tested: (C1) the trap persists unaided; (C2) an integrated portfolio beats single-stance policy; (C3) the connective-failure case responds to market linkage; (C4) the two archetypes diverge by which lever releases them; (C5) the integrated portfolio escapes. The decisive observation: the parameter that moves the income *level* most (the headroom term) barely touches the *conclusions* — the numbers are soft, the structural conclusions firm. Full detail is in the Sensitivity & Robustness Analysis.

14. Scope and limitations

- **A reasoning instrument, not a forecaster.** The model reports reliability and ranges, not point predictions; roughly 61% of parameters are declared assumptions.
- **Cases are illustrative.** The two archetype cases are demonstration instances, not a population sample or a validated calibration (see the Clarification Notice); their starting values are source-informed only.
- **Shared national thresholds.** The income lines are a single national per-capita set applied to both cases for comparability; state-specific lines (UP vs Telangana) would be more precise and are a natural refinement.
- **Structural validation only.** Validation covers structure and behaviour patterns; the model is not calibrated to, or tested against, empirical time series, and makes no out-of-sample predictive claim.
- **Reproducible but stochastic.** Results are exact for a given seed and vary across seeds; comparisons should hold the seed fixed.

15. Notes for reproduction

- One tick = one month; regime confirmation and fatigue are counted in months, the political/NPF update is annual (month 12).
- Income is a **level**, so to reproduce a trajectory it suffices to reproduce the `skills`, `market`, and `middlemen` stocks plus `subsistence`/`headroom`; the flow-style computation preceding the level assignment in `stepIncome` is largely vestigial and does not set income.
- Umbrella schemes contribute via their children (`UMBRELLA_CASCADE`), never double-counted.
- All constants above are the shipped defaults; the Admin panel and `updateCalibration()` can override any of them at runtime.

Cross-reference: the theory behind each mechanism is in the Theoretical Foundations Companion; the full test-by-test validation is in the Validation Report and the complete parameter sweep in the Sensitivity & Robustness Analysis. Abboudi, B. (2026). CPAF v6, Jindal School of Government and Public Policy, O.P. Jindal Global University.