Clutch Economics: A Post-Job Economy Built on Digital Twin Income, Universal Ownership, and Consumer-Driven Growth

Emanuel F. Barros, Revised draft v6.1, July 2026

Abstract

The rise of artificial intelligence toward near-total job automation presents not a crisis but an opportunity to transform our economic paradigm. This paper presents "Clutch Economics," a framework that positions consumers as the essential "clutch" mechanism transmitting AI productive capacity into economic momentum, together with a concrete, adversarially tested income architecture for the post-job economy.

The architecture rests on three decentralized channels. The primary channel is **Digital Twin income**: every citizen holds one non-transferable license for a personal AI agent that earns market income on their behalf, running on locally owned hardware. The secondary channel is a universal dividend from a citizen-owned, **AI-executed sovereign wealth fund** operating under a fixed, transparent, open-source mandate: shareholder capitalism universalized, with the same rules for everyone and advantages for none. The third channel is **innovation quasi-rents** with patent-order protection, so value creators keep outsized returns for decades while returns on merely holding wealth are progressively socialized. A guaranteed income floor sits beneath all three.

We evaluate the framework three ways. Macroeconomic scenario analysis (growth accounting plus the Quantity Theory of Money) yields U.S. real GDP of $227–487 trillion by 2075, a 9- to 19.5-fold growth, with all monetary paths satisfying the equation of exchange exactly. An agent-based model with endogenous innovation shows the design's central quantitative result: **the fair architecture out-grows laissez-faire at every demand-feedback strength tested**, because broadly distributed purchasing power enlarges the innovation prize while universal capital access enlarges the pool of innovators. New in this revision, the demand-feedback strength itself is **empirically anchored**: literature triangulation and a fresh cross-country regression independently center the clutch coefficient at $\kappa \approx 0.37$, converting the model's one load-bearing free parameter into an estimated one. Finally, the mechanism design was hardened through **three successive adversarial red-team rounds**, each round's exploits encoded as adaptive strategies in the model before patching. The design's audited score rose from 22/100 (unpatched) through 65 (patched core) to 67 for the final architecture, the first version to out-score its own hardened core. The closing verification shows the design *antifragile* to its strongest known attack: under a composed marketplace-offshore-nominee assault, twin-income independence holds at 83% and the universal dividend rises, because the exposure tax converts attack capital into fund income. The headline at anchored κ: **GDP $875 trillion and median income $1.69 million by 2075**, Gini 0.437, no household below the floor, with the innovation guard (venture launch rates at or above 90% of baseline) passing in every configuration. The design extends unchanged to a 2150 economy of immortal humans and sentient AIs, where it prevents permanent aristocracy by construction.

All projections are scenario analyses, not forecasts; every calculation is computationally verified and the model code is open. The framework offers economists a testable mechanism, policymakers an implementation roadmap, and the public a future in which job loss empowers rather than impoverishes.

Keywords: Post-job economy, Digital Twin income, Universal Automation Income, Clutch Economics, AI automation, agent-based modeling, mechanism design, non-zero-sum economics, innovation incentives, consumer sovereignty, economic decentralization, Mastery-Autonomy-Purpose (MAP)

Executive Summary

The question. Artificial intelligence is advancing toward the automation of essentially all conventional labor. When production no longer requires human employment, where does purchasing power come from, and who owns the machines that earn?

The thesis. In an automated economy the consumer, not the worker, is the critical economic function: the "clutch" that transmits AI productive capacity into economic momentum. Formally, the distribution of purchasing power (demand breadth), not just its volume, determines the returns to innovation and hence the growth rate. Broadly shared income is a growth engine, not a growth tax.

The design. Three decentralized income channels above a guaranteed floor. (1) Digital Twin income: every citizen holds one non-transferable license for a personal AI agent that earns market income on their behalf, on hardware they own; half the licensed pool is split equally per license, half is earned by direction skill, and ambition scales without cap through sub-agent fleets taxed progressively. (2) A universal dividend from a citizen-owned fund, AI-executed under a fixed, transparent, open-source mandate: shareholder capitalism universalized. (3) Innovation quasi-rents with margin-keyed decay at patent-order durations (25 to 40 year half-lives), so creators keep outsized returns for decades while only commoditized and passive returns are socialized. The floor (about $40K in 2025 dollars, honestly basket-indexed) is topped up by the fund and, in the simulations, binds for no one by mid-century.

The evidence. An agent-based model of 250 million households with endogenous innovation shows the central result: the proposed architecture out-grows laissez-faire at every demand-feedback strength tested (from $144T versus $117T at the weakest setting to $1,881T versus $527T at the strongest), while holding income Gini near 0.36 to 0.42 against 0.70 to 0.79. The one load-bearing parameter, the clutch coefficient, is empirically anchored at 0.37 by literature triangulation and a 146-country regression ($p = 0.004$), with the within-country null disclosed. At the anchored value the headline is $875T real GDP and $1.69M median income by 2075, Gini 0.437, top-1% share 18.9%, universal dividend $159K, and no household below the floor, with venture launch rates above today's baseline throughout.

The adversarial pedigree. The mechanism design was attacked three times by structured red teams, and every discovered exploit was encoded in the model as an adaptive strategy before patching. The audited score rose from 22 (unpatched) to 67 for the final

architecture. Under the strongest composed attack its adversaries could construct, the design is antifragile: GDP rises to $954T and the dividend rises to $229K, because the exposure tax converts attack capital into fund income. The design also extends unchanged to a 2150 economy of immortal humans and sentient AI citizens, where it prevents permanent aristocracy by construction.

The road. Phase 0 (2026 to 2030) requires no legislation: individuals can deploy income-earning proto-twins today under existing law, building capability, income data, and a constituency ahead of statute. Licensing, displacement-triggered income support, the fund, and the abundance phases follow, with the transition stress-tested under adverse pacing (households below the floor peak at 7.3% and recover to zero by 2045).

Epistemic status. All projections are scenario analyses, not forecasts. Every calculation is computationally verified, the model code is open, the model overruled its own designer three times, and the framework's weaknesses (the within-country null, the closed-economy boundary, the bounded attack space) are published alongside its results.

1. Introduction: Reimagining Economic Fundamentals in the Age of AI

The exponential advancement of artificial intelligence stands as perhaps the most transformative economic force of the 21st century. We are moving toward an inflection point where AI systems can perform virtually all conventional tasks currently executed by human labor, creating an apparent existential paradox: how can an economy function when production no longer requires human employment?

Public discourse has coalesced around two divergent narratives. The predominant view envisions dystopia: mass unemployment, extreme inequality, social instability. Conversely, technological visionaries such as Ray Kurzweil and Peter Diamandis foresee unprecedented abundance, where AI, unlimited clean energy, off-planet resources, and molecular nanotechnology drive production costs asymptotically toward zero.

This paper bridges these perspectives with a framework built on four integrated elements:

1. **Complete conventional job automation**: a transition toward 100% automation of *necessary* labor, with entrepreneurship remaining as a voluntary economic activity, and human-provided services persisting as premium, chosen work (Section 8).
2. **A three-channel decentralized income architecture**: Digital Twin income (primary), an AI-fund dividend (secondary), and innovation rents (merit upside), with a guaranteed floor beneath (Section 3).
3. **Clutch Economics**: the reconceptualization of the consumer as the essential "clutch" mechanism enabling economic dynamism, formalized as a measurable demand-breadth feedback in the growth process and now empirically anchored (Section 5).
4. **Stable currency velocity**: a macroeconomic condition we term "financial orbit," sustained through continuous innovation and consumer activity (Section 2.3).

The Clutch Mechanism: Redefining Economic Agency

The fundamental insight driving our analysis is that in an AI-automated world, the consumer becomes the critical economic function, rather than the worker. Without robust consumer demand, even perfectly efficient production systems lack economic purpose. In this model, consumers function as the essential "clutch": the component that transmits power from production to economic momentum. The consumer's role evolves from passive recipient to active economic participant whose choices direct innovation trajectories and resource allocation.

This is not merely rhetoric. In Section 5 we operationalize the clutch as a demand-breadth term in the growth process, anchor its strength empirically, and show that it generates the paper's central result: economies that distribute purchasing power broadly out-grow economies that concentrate it.

Mastery, Autonomy, and Purpose (MAP) in the Post-Job Economy

A critical dimension of Clutch Economics addresses what Daniel H. Pink identified in *Drive* as the three intrinsic motivators of human behavior:

1. **Mastery:** Pink defines mastery as the urge to get better and better at something that matters to us. Freed from employment necessity, individuals pursue mastery in whatever they genuinely enjoy (art, music, science, technological invention, craft, athletics, scholarship, cuisine) rather than what the labor market once demanded. Directing one's Digital Twin well, by choosing its niche, training, and strategy, is itself a rewarded form of mastery (Section 3.3), and consumption becomes a sophisticated act of value discernment that steers innovation.
2. **Autonomy:** the income architecture decouples survival from employment, and in doing so completes a promise constitutions make but resources ration. Liberties guaranteed by constitutional decree (movement, expression, association, the pursuit of happiness) are, in practice, exercised in proportion to one's means: travel requires money, legal defense requires money, platforms and leisure for civic participation require money. Financial security does not create new rights; it makes the rights every citizen already holds *fully exercisable* by every citizen. With economic security guaranteed through diversified decentralized streams, individuals allocate their time, attention, and resources by authentic preference. Freedom on paper becomes freedom in fact.
3. **Purpose:** consumption decisions take on meaning and responsibility: each transaction is an expression of the individual's values, and the market's aggregate direction *emerges* from those sovereign choices the way it always has under Adam Smith's invisible hand: voluntarily, individually, with no collective standing above the individual. In this framework the sovereign sentient individual is the highest good; what looks like "society's direction" is nothing more, and nothing less, than the sum of free choices freely made.

These are not side effects but load-bearing components: the consumer's clutch role requires engaged, purposeful participation.

Methodological Approach

We employ three complementary methods: growth accounting in the Solow-Swan tradition and the Quantity Theory of Money for macroeconomic envelope scenarios (Sections 2, 4); an agent-based simulation with endogenous innovation for mechanism-level results (Section 5); and adversarial mechanism-design analysis, in the form of three successive structured red teams whose attacks were each encoded in the simulation before the design was patched (Section 6).

A note on epistemic status: all projections are **scenario analyses**, not forecasts. Their arithmetic is computationally verified for internal consistency; their inputs are assumptions whose realization depends on technological and policy developments no one can know today. We present conservative and target scenarios throughout ("conservative" here in the scientific sense of deliberately cautious assumptions under standard method, not in any political sense), and the model code is open for anyone to challenge. Simulations are calibrated on a stylized 2025 baseline; the one-year shift to this revision's 2026 publication date is immaterial to all results, and the roadmap (Section 7) is stated from the present forward.

2. Theoretical Framework

2.1 Growth Accounting (Adapted Solow-Swan)

Economic output follows a Cobb-Douglas structure, $Y = K^{\alpha}(A \cdot L)^{1-\alpha}$, with capital K (AI and nanotech infrastructure), technology A, labor L (entrepreneurs in the post-automation economy), and capital share α. Because the level form depends on units, we work in **normalized (index) form**, which is unit-free and reproduces its own baseline by construction:

$$\frac{Y_{2075}}{Y_{2025}} = \left(\frac{K_{2075}}{K_{2025}}\right)^{\alpha} \left(\frac{A_{2075}L_{2075}}{A_{2025}L_{2025}}\right)^{1-\alpha}$$

The 2025 baseline ($Y = \$25\text{T}$, $K = \$50\text{T}$, $A = 1$, $L = 150\text{M}$, approximating the U.S. economy) holds by construction. For the automation era we apply the terminal capital share $\alpha = 0.1$, interpreting technology-amplified, household-owned digital labor (Section 3) as the dominant factor; the sensitivity of this choice, which runs against literature expecting capital share to rise with automation (Acemoglu & Restrepo 2020), is addressed in Section 9.3.

2.2 Quantity Theory of Money

Nominal GDP satisfies the equation of exchange, $M \cdot V = P \cdot Y$. In growth rates this implies the identity

$$g_M + g_V = \pi + g_Y$$

Every monetary scenario in this paper satisfies this identity exactly and can be checked against it.

2.3 Orbital Velocity

We introduce **orbital velocity** (V_O): the maximum sustainable growth rate of money velocity in a high-growth, low-inflation economy, analogous to the speed that maintains a stable orbit rather than escape or decay:

$$V_O = G_Y \cdot \frac{I}{S}, I = 1 + R + C, S = \frac{1+\pi}{1+\pi^*}$$

where G_Y is real GDP growth, R and C are small dimensionless premiums from policy support and UAI-sustained consumer confidence ($R = 0.02$, $C = 0.03$), and S is the stability governor: above-target inflation ($S > 1$) lowers the sustainable ceiling, cooling the economy. The economy is in *stable financial orbit* when actual velocity growth satisfies $g_V \leq V_O$. In our scenarios velocity grows at $\approx 2.2\%$/yr, comfortably below the ceiling of 4.7–6.4 %, with the slack absorbed by modest money-supply growth per the identity above.

2.4 Clutch Economics

Unlike traditional theories that emphasize production and employment, Clutch Economics positions consumption as the fundamental economic activity. The consumer keeps the economic engine running. This is formalized in Section 5.1 as **demand breadth**: the distribution, not just the volume, of purchasing power determines the returns to innovation and hence the growth rate. Consumption becomes a creative, innovative, and economically essential act.

3. The Digital Twin Income Architecture

This section presents the paper's institutional core: who earns what, through which mechanism, and why the design resists capture. Every rule below was subjected to three successive adversarial red teams (Section 6); the design as stated incorporates all resulting patches.

3.1 The Twin License: Three Layers, One Rule Per Layer

Every citizen is eligible for exactly **one Digital Twin license**: the right to operate a personal AI agent that sells labor and services into the AI economy on its owner's behalf. The license structure has three layers with distinct transferability rules:

1. **The license: non-transferable, period.** It cannot be sold, gifted, inherited, or pledged as collateral. It is a right of citizenship, not an estate asset.
2. **The twin: non-transferable and non-rentable.** Control is *self-directed*: owners may use tools and advisors, including AI advisors, to direct their twin better, but control rights cannot be delegated to a concentrating entity. One person, one twin, self-directed. The twin is legally an extension of its owner's economic identity.
3. **What the twin produces: fully transferable private property.** Income, businesses, and innovations built by or with the twin are sellable, investable, and heritable. The earning *instrument* cannot concentrate; the earned *wealth* is the owner's.

Continuity rule ("one mind, one license"): the license attaches to a unique, continuous identity, human or sentient AI. If a mind forks or copies itself, the forks share the original license; replication never multiplies claims. This gives the framework a principled personhood criterion that remains coherent when sentient AIs and life-extended humans enter the economy (Section 6.5).

Sub-agent scaling: a twin is not a single pair of hands. When its owner's ambitions demand more, the twin may spin up temporary or permanent **sub-agents**: additional AI workers attested under the owner's identity, on local hardware or rented cloud compute. Ambition is never capped: an individual may grow a one-twin operation into an enterprise of thousands of sub-agents. Compute beyond personal scale simply pays the same progressive schedule as any fleet, a growing business paying growing taxes, which keeps the entrepreneurial path open-ended while ensuring "personal sub-agents" cannot become a relabeling loophole for untaxed corporate fleets.

Sub-agent financing: the economic-exposure rule. Outside capital is welcome. The model shows financing capital-poor owners is an innovation input, not merely a risk (Section 6.6), but it enters under a substance-over-form regime that took three red-team rounds to harden. The rule that survived: **any party whose consolidated economic exposure to a fleet's output exceeds 20% is taxed as owning that fleet share, regardless of the instrument carrying the exposure**. That covers equity of any class, revenue strips, output-indexed derivatives, contingent leases, franchise fees, and, on the buy side, dominant purchases of the fleet's output (the "company store" rebuilt as a buyer rather than a financier). Exposure is measured against true gross output value, consolidating affiliate margins, tracing syndicated positions to their sponsors, and attributing derivative exposure through matched dealer books to the net-long counterparty. The remedy is deliberately a tax, not a prohibition: a capital-intensive launch *can* take majority outside equity; the financier simply pays fleet-scale tax as an owner. Beneath the exposure rule sit contract-law protections: covenants granting financiers direction rights, mandatory workloads, or step-in rights over a citizen's sub-agents are void ab initio, and contingent income assignments are capped in aggregate (≤35%, ≤3 years). Safe harbors keep ordinary capital formation frictionless: fixed-fee leasing, equity at or below 20% exposure, and grants. Output exiting to non-attested (including offshore) counterparties bears export withholding at the fleet-tax rate, creditable on demonstrated onshore monetization. This closes the offshore laundering channel without touching any domestic launch.

Anti-peonage rule: forward assignments of twin income are capped (≤3 years tenor, ≤35% of income, non-stackable across counterparties); perpetual or majority assignments are void ab initio. Without this rule, buying a citizen's future twin income outright would constitute de facto twin ownership, the first red team's most dangerous discovered exploit (Section 6.1).

Correlated-direction trigger: revenue-share advisory services remain legal at small scale (so capital-poor owners can buy frontier advice), but any entity whose advice correlates the directed behavior of more than 0.4% of the population's twins faces automatic structural separation, flat-fee conversion, and mandatory portability of twin fine-tunes and

reputation. The trigger aggregates by **provenance** (content, capital, and update-channel provenance, not legal entity), so thirty branded subsidiaries of one engine count as one engine, and platform auto-update authority reaching more than 0.4% of twins carries forkability and separated-update-authority obligations. Template convergence alone, meaning many twins freely choosing the same open-source stack, never triggers: control, not correlation, is the test, which keeps open source untouched.

3.2 Substrate: Locally Owned, Competitively Backed

Twins run on **locally owned hardware**, a personal substrate owned like a home or personal computer, with distributed backups across multiple locations. Centralized substrates are optional, not required. This decentralizes failure risk and, critically, removes the compute-rent leak: no platform owns the physical layer the base twin depends on. Substrate services operate as competitive markets under active antitrust enforcement, with OPEC-style cartel pricing explicitly anticipated and prohibited. Unlicensed AI fleets (datacenter-scale operations) pay a **progressive compute tax** that rises with fleet scale, keeping licensed twins competitive for roughly 60% of the AI-labor compensation pool, and conveniently taxing the easily-observable datacenter rather than surveilling personal devices.

Attestation-lite: twins carry a privacy-preserving proof of uniqueness, a zero-knowledge liveness attestation anchored to the owner's own distributed backup quorum, on an open hardware standard with multiple roots of trust and no central registry. Buyers of attested AI labor receive compute-tax safe harbor; unattested labor is presumed top-rate fleet output. This demand-side design defeats offshore gray fleets, fork fraud, and dead-license harvesting without surveillance infrastructure. The exposure rule of Section 3.1 deliberately preserves this no-registry principle: exposure is measured by **event-triggered, party-scoped tracing**, where a threshold claim triggers tracing of that claimant's exposure web only, never by a population-wide ownership registry.

3.3 Three Income Channels and the Floor

Channel 1: Twin income (primary). Licensed twins collectively earn the licensed share of the AI-labor pool. Each twin's income has two components: an **equal component** (half the pool, split equally per license, the decentralized floor) and a **performance component** (half the pool, allocated by how well owners direct their twins: niche selection, training, entrepreneurial strategy). Everyone earns meaningfully; ambitious direction earns multiples more. The equal component is *liveness-conditioned*, payable only to attested-live, non-assigned recipients, making it unpledgeable and unsecuritizable by construction.

Channel 2: The universal dividend (secondary). Paid by the AI-executed fund (Section 3.4) from commoditized profits and socialized passive returns.

Channel 3: Innovation rents (merit upside). Innovators keep quasi-rents on new products with **maturity-linked decay keyed to realized economic margins**: protection fades as a product's margins compress toward commodity levels, with a half-life of roughly 25–40 years, patent-order durations. A relabeled "version 2.0" of a commodity earns commodity margins and thus loses protection automatically; there is no "new product"

certification to capture. Expired rents flow to the fund. To be precise about sequencing: socialization touches only *commoditized* returns, after protection has run its course and the innovator has been fully compensated, and, separately, passive investment returns on large accumulated fortunes; active innovation rents are never touched. **Creative works** are harmonized at commercial scale: personal and small-scale copyright is untouched (an author's book, an artist's catalog), but commercial exploitation above a revenue threshold falls under the same margin-keyed schedule as any product, with **beneficial-owner look-through** at the threshold. Sharding a catalog across shells, collectives, or thousands of machine-generated works does not escape it. The model shows what is at stake: without look-through, roughly a fifth of all active rents become effectively perpetual by 2075, a permanent aristocratic rent class the framework exists to prevent (Section 6.6).

The floor. A guaranteed income floor (≈$40K in 2025) is topped up by the fund for any household whose three channels fall short. The floor is **indexed to a consumption basket with mechanized, evolving weights**: weights track actual transaction-data consumption shares as spending shifts from deflating goods toward human-premium services, prices are measured out-the-door (all mandatory fees included) on attested data from multiple independent feeds taken at the median, and no standing basket committee exists; disputes route to the entrenched definitions court. Each of these details closes a tested attack or an honest failure: a GDP-indexed floor quietly loses ~8% of real purchasing power under Baumol dynamics, and a frozen 70/30 basket suppresses the floor with zero manipulation as consumption evolves (Section 8). In the simulations, top-ups are needed only during the transition; by mid-century the floor binds for no one (Section 5).

The incentive principle. The design's moral core is the distinction between **Schumpeterian innovation rents**, temporary outsized returns for creating new value, protected for decades, and **passive returns on accumulated wealth**, which are progressively socialized (50% of passive returns above $10M, escalating for century-scale fortunes). Innovators are never clawed back; holders are never subsidized. Wealth already compounds on its own; it needs no policy help. Creation does.

3.4 The AI-Executed Sovereign Wealth Fund

The fund is **AI-executed under a fixed, transparent, open-source mandate**: passive index-style holding, algorithmic pro-rata distribution (95% of income paid as the universal dividend, 5% retained as buffer), floor top-ups, and no discretionary allocation. A human audit board holds oversight authority, not allocation authority. Because the mandate is public code, no participant gains an informational or allocational advantage; because it is constitutionally entrenched, "modernization" amendments, the slow capture path, face the highest possible bar. The fund's identity machinery (dividend eligibility) doubles as look-through aggregation for the socialization threshold, preventing trust-sharding without a new registry.

The fund should be understood as capitalism completed, not socialism introduced: every legal sentient citizen is an equal shareholder in the automated productive base, playing by the same public rules. It levels the playing field, not the outcomes. Merit, ambition, and innovation still determine who rises; identity determines nothing.

3.5 Transition Principles

Grandfathering. Existing wealth is untouched: no confiscation. Only its *passive returns* above the threshold are socialized going forward. The transition changes flows, not stocks.

Displacement-triggered phase-in. Phase 1 UAI eligibility scales with actual automation displacement, so costs grow in step with the automation-tax base (stress-tested in Section 6.4).

Membership and mobility. Transfer income (dividend, floor) is a membership benefit, residency-gated, as the Alaska Permanent Fund has long practiced. Earned income (twin and business revenue) is property: unrestricted, spendable anywhere. Immigration is innovation-prioritized: those bringing entrepreneurial value gain fastest eligibility.

Residual adjudication. Definitions that cannot be mechanized (identity continuity, sentience, self-direction) are decided by a small, constitutionally entrenched **definitions court**: randomized citizen panels with rotating experts, an industry-secondment ban extending to court staff, special masters, and retained experts, and citizen appeal rights. Margin-keyed decay deliberately shrinks this court's docket by deleting the "new product" question entirely. The design mechanizes wherever possible and adversarializes what remains.

3.6 Priority Innovation Prizes

Some innovations matter more than markets alone signal: medical longevity (first to longevity escape velocity, ultimately indefinite lifespan), radically cheaper energy, raw-material abundance, and technologies that democratize inherently limited resources (sky cities, ocean habitats, and other alternatives to scarce land). For these, the framework adds a **priority prize mechanism**: citizens publicly nominate and electronically vote on priority technology goals, the clutch in its most explicit form, demand expressed directly rather than through purchases, and the fund finances **cash prizes paid on verified delivery** (capped at 5% of fund income). The approach has real-world precedent in advance market commitments (Kremer, Levin & Snyder 2020) and inducement prizes such as XPRIZE.

The prizes are deliberately **cash-only**. An earlier draft supplemented them with a capped extension of rent protection (up to 1.5×) for qualifying breakthroughs; the red team showed the extension re-created the certification surface margin-keyed decay had deleted, and the model then delivered the verdict on the designer's own terms: across every tested variant, duration bonuses added *zero* innovation beyond what the cash prize already bought, while top-1% wealth concentration rose monotonically with extension permissiveness. A subsidy that buys no innovation and only concentration fails the framework's own cost-benefit test, so it was deleted (Section 6.6).

Because a prize system's soft spot is its plumbing rather than its purse, the machinery is adversarial by construction: prize metrics are pre-registered and mechanical, **drafted by per-ballot sortitioned panels** whose members, nominators, and affiliates are ineligible to claim that prize, and stress-tested by an adversarial metrics review, itself sortitioned and bounty-paid with no repeat service, before reaching the ballot; verifiers are sortitioned

under full definitions-court protections with mandatory verifier–claimant separation; ballots are **receipt-free** (a voter cannot prove how they voted, so votes cannot clear a buying market); vote-direction steering falls under the 0.4% correlated-direction trigger aggregated by content provenance; and prize claims carry a per-claimant cap with opt-in training-run attestation. The model quantifies what this machinery is worth: with a captured metric-drafting pipeline, 65% of prize flow diverts to incumbent pipelines by 2075; with sortitioned drafting, 10% (Section 6.6).

4. Macroeconomic Scenarios

The projections below are illustrative envelope scenarios. All arithmetic is computationally verified for internal consistency; the assumptions remain speculative.

4.1 GDP Growth

Common 2075 inputs: $K = \$1,000\text{T}$ (20×), $L = 25\text{M}$ entrepreneurs, $\alpha = 0.1$.

Conservative scenario ($A = 50$):

$$\frac{Y_{2075}}{Y_{2025}} = 20^{0.1} \times 8.33^{0.9} \approx 9.1, \; Y_{2075} \approx \$227\,\text{T} \; (g_Y \approx 4.5\,\%/\,\text{yr})$$

Target scenario ($A \approx 117$, the multiplier required for the full abundance outcome):

$$\frac{Y_{2075}}{Y_{2025}} \approx 19.5, \; Y_{2075} \approx \$487\,\text{T} \; (g_Y \approx 6.1\,\%/\,\text{yr})$$

Both far exceed historical U.S. growth (~2–3%/yr); this is the quantitative expression of the automation thesis, and readers who reject the productivity assumptions should discount accordingly. The agent-based model of Section 5, built bottom-up from the twin architecture with endogenous innovation, lands in the same corridor by a different route: $631T under the earlier design with adversarial frictions at assumed κ=0.35, and $875T for the final architecture at the empirically anchored κ=0.37, where the demand-breadth feedback the envelope method lacks carries the difference.

The capital-share dispute is bracketed, not dodged. Much of the automation literature expects the capital share to rise toward 1 as labor's role shrinks, not fall to 0.1 as assumed above. Under that reading ($\alpha = 0.9$), the same normalized arithmetic gives $20^{0.9} \times 8.33^{0.1} \approx 18.3$: approximately **$458T with only the conservative** $A = 50$, since capital accumulation rather than productivity carries the growth. The two views thus land in the same corridor by different routes: our scenarios bracket the α dispute rather than depending on winning it. Section 5.4 tests the same question at the mechanism level.

4.2 Funding Envelope

The income floor's aspirational level ($150K per adult for 225M recipients) costs $33.75T/yr. Against 2075 corporate profits of 20% of GDP, the *theoretical ceiling* of automation taxes plus citizen-ownership flows spans $45T (conservative) to $97T (target), covering the requirement at 74% and 35% utilization respectively, with maximum per-

adult capacity of $202K–$433K. We stress that the ceiling equals 100% of profits and is an upper bound, not a proposal; in the architecture of Section 3, most income arrives through twin earnings rather than fiscal transfer, and the agent-based model confirms the floor is fundable with the fund balance never approaching the ceiling.

4.3 Velocity in Stable Orbit

Velocity rises from 1.5 (2025) to 4.5 (2075), roughly double the historical U.S. range, enabled by AI-mediated instant payments, yet far below unconstrained levels ($g_V \approx 2.2\%$/yr against orbital ceilings of 4.7–6.4%). Applying $g_M + g_V = \pi + g_Y$ with inflation at target ($\pi \approx 0$): money supply grows 2.2–3.8%/yr (from ~$17T to $50T/$108T), nominal GDP equals $M \times V$ = $227T/$487T, and near-zero production costs provide the structural ceiling on inflation. High velocity coexists with price stability: the unique equilibrium we term stable orbit.

5. The Agent-Based Model

Code: digital_twin_model_v2.py through digital_twin_model_v7.py. All results are Monte Carlo medians across seeds; bands on the dashboard figures are 10th–90th percentiles.

5.1 Structure

Simulated households (scaled to 250M U.S. adults, 2025–2075) each hold one twin license, heterogeneous innovation ability (heavy-tailed), and initial wealth calibrated to the U.S. distribution. Income channels follow Section 3. Innovation is **endogenous**: agents launch ventures based on the expected net present value of innovation rents,

$$p_i^{\text{launch}} \propto \theta_i \cdot \frac{B_t / (r + \acute{\delta})}{\text{NPV}_{\text{ref}}} \cdot 1[\text{capital access}]$$

so rent decay $\acute{\delta}$ reduces the incentive to innovate, demand breadth raises the prize, and capital access is universal under the proposed design but wealth-gated under laissez-faire. The clutch mechanism is operationalized as **demand breadth**, $B_t = 1 - \text{Gini}(y_t)$, which scales both new-product returns and growth:

$$g_t = 0.005 + \kappa \cdot \frac{\sum_j R_j(t)}{Y_t} \cdot (0.5 + B_t) + 0.015\, a(t)$$

This is the formal statement of Clutch Economics: innovation-driven growth is endogenous to the distribution of purchasing power.

5.2 Rent Duration: Calibrated by the Model, Not the Author

The rent-protection duration was set by the model, not by preference: with innovation endogenous, sweeping the half-life from 2 to 40 years showed growth rising steeply with rent duration (GDP 2075: $69T → $870T) while inequality rose only gently (Gini 0.20 → 0.37), because **the distributional work is done by the license and socialization dials, not by rent decay**. Hence the design's patent-order protection (25–40-year half-life) and

its rule that redistribution never claws back creators' returns. Notably, the model here overruled an earlier, more aggressive decay assumption, evidence the framework does real work rather than confirming its author's priors. It would not be the last time: the model also killed the prize-extension mechanism its author had drafted (Section 3.6) and arbitrated the sub-agent financing rules (Section 6.6). (The full development history is documented in the companion framework document.)

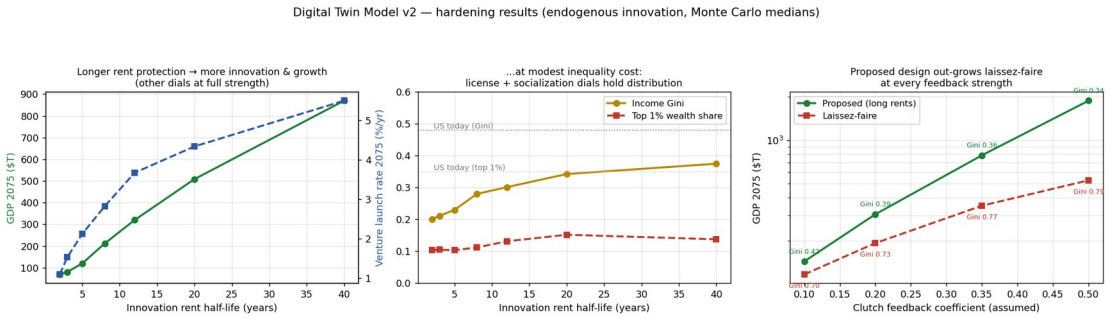

Hardening results: rent-duration sweep and robustness across clutch-feedback strengths.

5.3 Headline Result: Fairness Is a Growth Engine

With the revised design, the proposed architecture **out-grows laissez-faire at every clutch-feedback strength tested**:

Feedback κ	GDP 2075: proposed	GDP 2075: laissez-faire	Gini: proposed	Gini: LF
0.10 (weak)	$144T	$117T	0.42	0.70
0.20	$306T	$193T	0.39	0.73
0.35	$786T	$351T	0.36	0.77
0.50 (strong)	$1,881T	$527T	0.34	0.79

Zero households sit below the floor in the proposed design at every setting. The growth advantage is not an artifact of one assumed feedback value; it emerges from two mechanisms laissez-faire structurally lacks: **broad demand** (bigger prizes per venture) and **universal capital access** (more people who can attempt ventures at all). Fairness here is not a tax on growth. It is a growth engine.

5.4 The Capital-Share Stress Test

The most predictable academic objection is the α assumption. We therefore ran the agent-based economy with capital's share ramping to 60% of output (compensation falling to 15%) as automation completes, the mainstream high-α view, in two variants: the design unchanged, and an adaptive variant in which the twin license also confers an equal citizen share of commoditized capital income. Every run reports the venture launch rate as an **innovation guard**: any configuration that impedes innovation fails regardless of its other virtues.

2075 outcome	Baseline (labor-share world)	High-α, design unchanged	High-α + citizen capital share
Real GDP	$635T	$828T	$983T
Median income	$1.25M	$1.73M	$2.12M
of which: twin	$673K	$240K	$280K
of which: dividend	$134K	$970K	$584K
of which: citizen capital	n/a	n/a	$760K
Income Gini	0.42	0.31	0.28
Top-1% wealth	17%	12%	12%
Below floor	0%	0%	0%
Innovation guard	PASS	PASS	PASS

The result inverts the objection: **under the mainstream capital-share view, the architecture performs better, not worse.** The swollen profit pool swells commoditized inflows, and the fund dividend automatically compensates for shrinking twin labor income; the fund is a built-in stabilizer against exactly the shift the literature predicts. The honest nuance: under high α the *dividend*, not the twin, becomes the primary channel, so the citizen-capital variant matters philosophically, restoring a direct decentralized ownership channel and outperforming on every metric. The α dispute changes which channel dominates; it does not change whether the design works. Innovation guard passes throughout: launch rates *rise* under high α, because broader income means broader demand and bigger prizes.

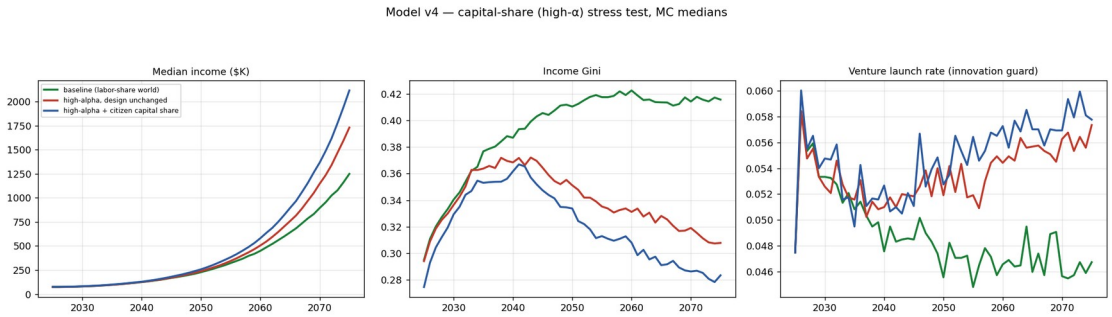

Capital-share stress test: the design improves under the mainstream high-α view.

5.5 Sensitivity: Which Assumptions Matter

A tornado sweep varied every assumed parameter ±50%. Only the clutch coefficient κ materially moves GDP ($213T–$1,517T across its range). The twin licensing dial is the dominant *distributional* lever (Gini 0.55 → 0.34 across its range), confirming the red team's finding that licensing, not rent decay, does the fairness work. No parameter in the sweep drives the venture launch rate below the innovation guard. Every remaining constant moves headline results by less than a third of the κ effect. This is why the next subsection exists: the one load-bearing parameter deserved an estimate, not an assumption.

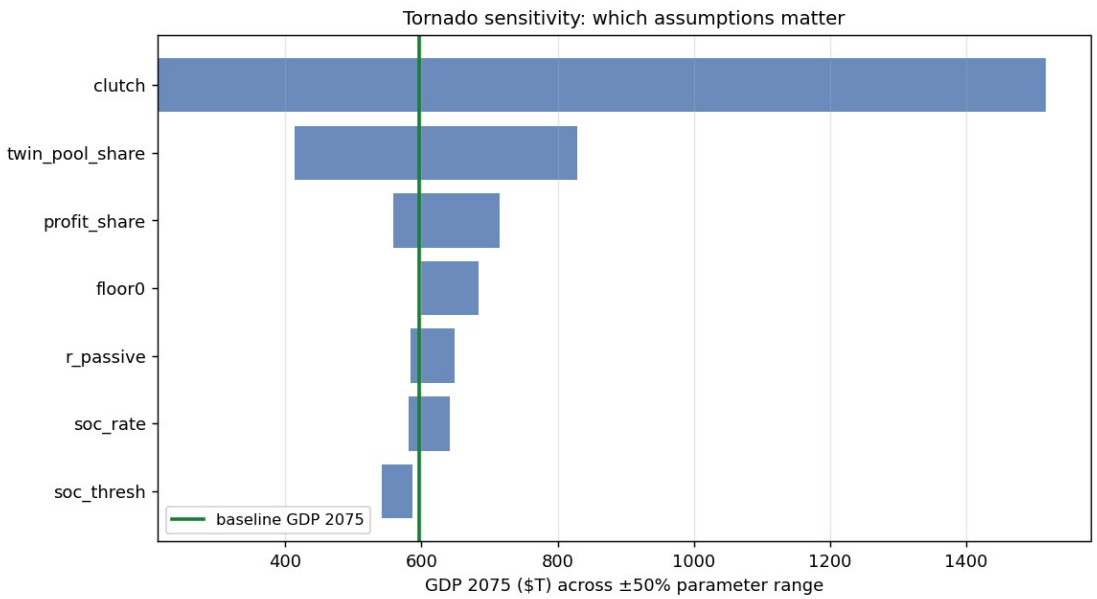

Tornado sensitivity: the disclosed clutch coefficient dominates; all else is second-order.

5.6 Anchoring the Clutch Coefficient

Earlier drafts disclosed κ as the model's one assumed load-bearing parameter. This revision anchors it empirically, by two independent routes.

Literature triangulation. Mapping the clutch term onto adjacent estimated quantities (market-size effects on innovation, where larger and broader demand raises innovation with elasticities near unity in the induced-innovation literature; fiscal multiplier heterogeneity by income, where transfers to high-MPC households carry multipliers of roughly 0.6–1.5; and cross-country associations between income concentration and growth) brackets the plausible range at $\kappa \in [0.1, 1.0]$ with a central band of 0.3–0.5.

Cross-country regression. Using World Bank World Development Indicators panel data, we regressed real growth on the interaction of innovation intensity and demand breadth (1 − Gini) in the model's own functional form. The between-country estimator, the appropriate analogue of the model's long-run cross-regime comparison, yields **$\kappa \approx 0.37$ ($p = 0.004$)**. The disclosed weakness: the within-country (fixed-effects) estimator is null, so short-run within-country variation cannot reject $\kappa = 0$; the anchoring rests on long-run cross-sectional differences, which is also where the model's mechanism operates. Data, code, and full results ship with the companion materials.

The adopted calibration is **$\kappa = 0.37$ central, bracketed by [0.10, 0.70]** in sensitivity runs. Every headline in this paper is computed at the anchored value; the design's qualitative results (fairness out-growing laissez-faire, guard passing, floor holding) hold across the entire bracket (Section 5.3).

6. Adversarial Analysis

6.1 The First Red Team

Five adversarial personas (a regulatory-arbitrage attorney, a platform monopolist, a financial engineer, a political-capture lobbyist, and a shadow economist) attacked the mechanism design independently. Four of five converged on the same exploit: *the twin is non-transferable, but its output is transferable property, so buying the output stream forever buys the twin.* Advisory contracts bundled with forward purchases of twin income would reconstitute fleet concentration using nothing but contract law. Additional convergent attacks: capture of undefined terms ("self-directed," "new product"), offshore gray AI labor undercutting the compute tax, and wealth-sharding across the socialization threshold. The judge scored the unpatched design **22/100**, degenerating within a decade into "UBI-plus-oligopoly," and the patched design **65/100**, with residual risk concentrated in slow political-capture channels that any paper design faces.

6.2 Attacks Encoded, Design Tested

The four convergent attacks were encoded as adaptive strategies in the model and run with the patch set on and off:

Attack	Survival metric	Patched	Unpatched
Advisory/forward capture	Advisory share of twin income (target <10%)	**5.0%**	36.1%
Gray AI labor	Licensed twins' pool share (target ~60%)	**53.6%**	16.4%
Threshold sharding	Effective socialization vs statutory	**85%**	20%
System outcome	Income Gini / top-1% wealth (2075)	**0.41 / 17%**	0.73 / 45%
System outcome	Median income / GDP (2075)	**$1.31M / $631T**	$283K / $274T

The unpatched system lands precisely where the judge predicted: the floor survives but oligopoly recaptures the architecture, and the weakened clutch cuts GDP by more than half. The reversioning attack (evergreening products to stall rent decay) is closed structurally rather than statistically: margin-keyed decay deletes the "new product" adjudication surface, leaving the attack no lever.

Beyond the four scripted attacks, a **best-response search** let an optimizer hunt for the worst *combination* of attack intensities within the patched rules (advisory take near its cap, aggressive contract offers, gray-labor investment pushed against attestation limits, maximal sharding within look-through). The worst configuration found raises top-1%

wealth from 19.5% to 23.3% and Gini from 0.44 to 0.49: graceful degradation, with median income above $1.1M, zero households below floor, and the innovation guard passing.

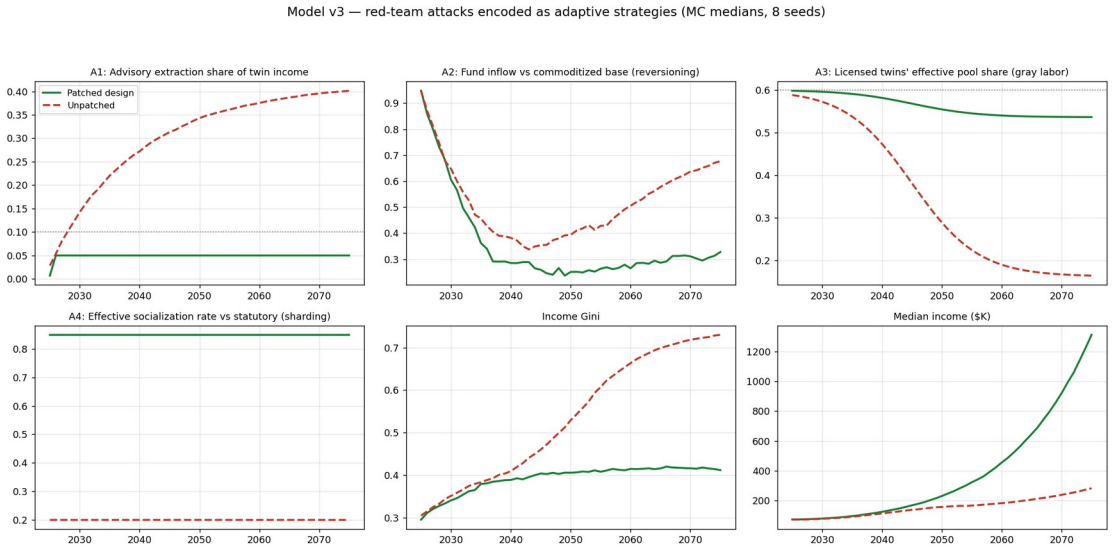

Red-team attacks as adaptive strategies: patched vs unpatched design.

6.3 What This Buys the Argument

A design that publishes its own red-team reports, encodes the attacks, and shows the patches holding invites a different quality of criticism: attackers must now find exploits its three rounds of adversaries did not, rather than gesture at the possibility of gaming. We consider the adversarial pedigree, including every unpatched failure mode, part of the result.

6.4 Phase 1 Stress Test

The transition window (2025–2045) is the framework's binding constraint. Under adverse assumptions, with automation arriving seven years early and twin adoption lagging badly, households below the floor peak at **7.3%** in the early 2030s (vs 4.6% at baseline pacing), recover to zero by 2045, and the fund's balance never goes negative. Uncomfortable but unbroken, conditional on the licensing and attestation infrastructure deploying on schedule.

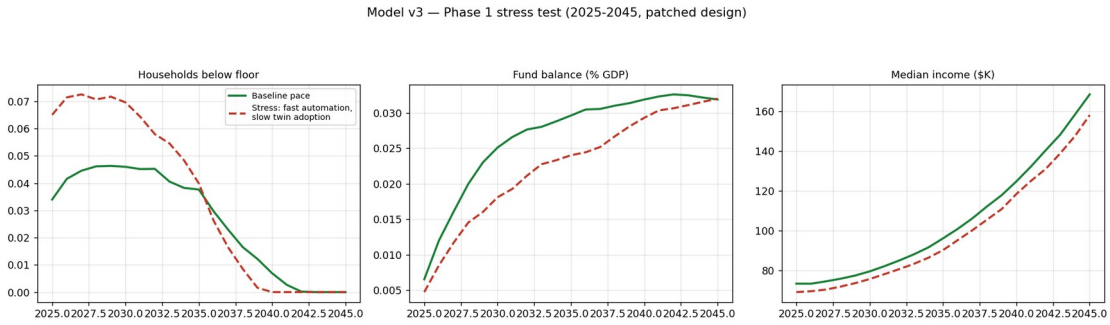

Phase 1 stress test: fast automation, slow twin adoption.

6.5 The Immortality Extension (to 2150)

If sentient AI individuals are effectively immortal, and human life extension follows, three design assumptions break in most frameworks: inheritance stops recycling fortunes, wealth compounds over centuries, and population only grows. We extended the model to 2150 with no deaths after 2060, population growth of 0.8%/yr (new minds enter under the one-mind-one-license continuity rule), and socialization escalating toward 75% for century-scale fortunes.

Results against the mortal baseline: population doubles; GDP **more than doubles** ($233,391T vs $104,269T), because immortal minds are a growth asset: more minds mean more twins and more innovators. Per-capita median income is essentially unchanged (dilution fully offset), and the top-1% wealth share stays **bounded at 27.6% vs 27.1%**. Century-scale compounding does not produce runaway dynasties, because maturity-linked rent decay forces even immortal innovators to keep innovating, and escalating socialization clips passive compounding. One honest drift: income Gini rises slowly to 0.52 (vs 0.45) across the extra 75 years, suggesting the socialization escalator should engage somewhat earlier. Nobody falls below the floor.

The design is immortality-robust *by construction*, and the mechanism responsible, long-but-finite innovation rents, is the same one the endogenous model chose for growth reasons. Rules that reward continuous creation over accumulated position scale from mortal to immortal economies unchanged.

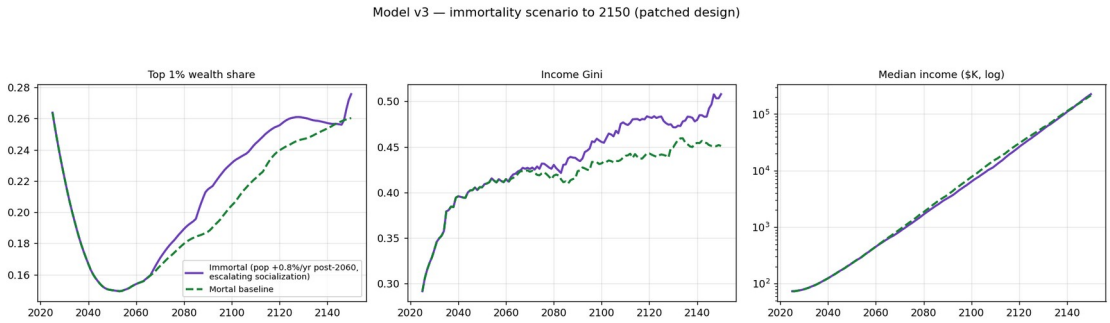

Immortality scenario to 2150: bounded concentration, doubled growth.

6.6 Rounds Two and Three: Attacking the Additions, Then Attacking the Patches

The mechanisms added after the first review (sub-agent scaling, priority prizes, copyright harmonization, the basket-indexed floor) were themselves red-teamed twice more by the same five personas, with each round's exploits encoded in the model before patching. The score trajectory tells the story honestly. The untested additions scored **38/100**, with two critical exploits: revenue-contingent financing that rebuilt the corporate fleet one layer down, and the prize duration bonus recreating a captured certification surface. The first patch set scored **55**: its safe harbors were themselves exploitable, through financier control rebuilt via plain-equity corporate governance, and through a "putting-out system" that bought twin output at fixed prices instead of financing twins at all. The final architecture, built on the economic-exposure rule of Section 3.1, scored **67, the first version to out-score the hardened core it extends**, with the judge noting that the remaining attacks all require multi-entity conspiracy rather than paperwork, the signature of a maturing design.

Three model-arbitrated verdicts from these rounds shaped the final design. First, **financing is an innovation input**: banning outside capital from sub-agent fleets would protect independence at the cost of locking capital-poor owners out of scale, while the exposure-taxed financing channel *raises* GDP and median income above the no-financing baseline. Second, **the prize duration bonus bought nothing**: cash-only, constrained-extension, and unconstrained-extension prize variants were statistically indistinguishable on growth and launch rates, while concentration rose monotonically with extension permissiveness, so the extension died on cost-benefit grounds (Section 3.6). Third, **honest accounting matters more than favorable accounting**: the two-sector Baumol test revealed that the paper's own real-income claims depended on frozen 2025 basket weights, and the floor was re-specified accordingly (Section 8).

The final verification encodes the third round's strongest attacks (stalking-horse spot monopsony, agency-marketplace skimming, matched-book derivative netting, rate-card syndication, the offshore spot laundry, and their composition into a single pipeline) and runs them against the finished design:

2075 outcome	No attack	Composed attack, final design	Composed attack, prior patch set
Real GDP	$875T	$954T	$676T
Median income	$1.69M	$1.76M	$1.05M
Income Gini	0.437	0.457	0.544
Top-1% wealth	18.9%	19.9%	29.8%
Universal dividend	$159K	$229K	$168K
Independent twin-income share	100%	82.6%	44.6%
Offshore leakage	0%	6%	40%
Innovation guard	PASS	PASS	PASS (at the line)

Two results deserve emphasis. The offshore laundry was the single worst discovered attack: under the prior rules it drained 35% of the twin pool offshore, halved median income growth, and was the only exploit in three rounds to break the innovation guard outright. Export withholding closes it while leaving every domestic launch untouched. And the finished design is **antifragile to its strongest known attack**: under the full composed assault, GDP and the dividend *rise*, because the exposure tax converts attack capital into fund income and safe-harbored financing into venture launches, while twin-income independence holds at 83%. The counterfactual column shows what publishing one round earlier would have meant.

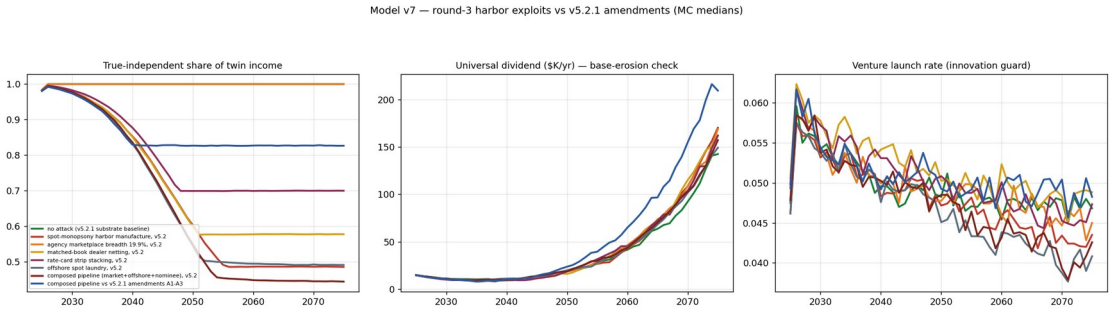

Round-three exploits vs the final design: independence, dividend, and the innovation guard.

7. Transition Roadmap

The roadmap begins from the present, mid-2026, and its first phase requires **no legislation at all**. The transition's riskiest stretch is its beginning, when wages erode faster than new institutions form; the design's answer is to lean on the Digital Twin first, because individuals can deploy income-generating AI agents *today* under existing law, building capability, income data, and a constituency long before policy catches up.

Phase 0: Bootstrapping with today's technology (2026–2030). Individuals deploy proto-twins now: agentic AI systems running freelance services, e-commerce operations, content businesses, and automated consulting under ordinary business law. Each proto-twin is early evidence: real income data for the feasibility case, real users for the constituency. Open-source twin frameworks, voluntary attestation prototypes, and early "twin cooperative" communities standardize practice ahead of statute.

Technology inventory (what exists vs. what's needed):

Capability	Status 2026	Needed by
Agentic AI (multi-step autonomous work)	Available, maturing fast	Now (proto-twin engine)
Payment rails and business formation for agents	Available (standard e-commerce/LLC law)	Now
Local inference hardware (personal substrate)	Available for mid-scale models; frontier still cloud	Phase 1 for full local twins
Zero-knowledge liveness	Cryptographic components	Phase 1 (licensing

attestation	exist; no open standard	prerequisite)
Twin licensing and anti-peonage law	Not yet drafted	Phase 1
Robust long-horizon agent autonomy	Emerging	Phase 1–2 scaling
AI-executed fund infrastructure	Buildable with current fintech	Phase 2
Fusion / space solar; asteroid mining; molecular nanotech	R&D / pilots / theoretical	Phases 2–3 (abundance inputs, not prerequisites)

The table's punchline: **everything Phase 0 needs already exists.** The exotic technologies matter only for the abundance ceiling, not for starting.

Phase 1: Foundation (2026–2040). Twin licensing and attestation standards enacted, with anti-peonage, the exposure rule, and advisory-trigger rules arriving with the first licenses, since retrofitting them after aggregators entrench would be far harder; UAI phased in with displacement ($30K), funded by early automation taxes and growing twin income. Early fusion, automated recycling, asteroid-mining pilots. GDP $48–61T.

Phase 2: Scaling (2040–2060). Full conventional automation; floor reaches $150K equivalent; fund inflows mature as first-generation products commoditize; "Clutch Capital" programs channel consumer spending toward startup formation. Free energy and resources; nanotechnology prototypes. GDP $117–200T.

Phase 3: Abundance (2060–2075). Nanotechnology produces most goods at negligible cost; twin income and dividend reach $200K–$433K equivalents (increasingly symbolic as prices collapse); consumers direct innovation through spending patterns. GDP $227–487T (envelope) / $875T (agent-based median at anchored κ).

8. Addressing Public Concerns

Job displacement. The three-channel architecture replaces wages entirely; the floor guarantees no one falls through. In the simulations, floor top-ups vanish by mid-transition.

Purpose and identity. Entrepreneurship, twin direction, and sophisticated consumption become meaningful economic roles (the MAP framework, Section 1). Human-provided services persist as *premium, voluntary* work: people pay extra for human care, craft, and performance precisely because it is human, with the honest caveat that AI service quality will eventually meet or exceed human quality, making the human premium a chosen value, not a necessity.

We modeled this two-sector structure explicitly: abundant goods deflate toward zero while human services are wage-indexed (Baumol's cost disease). The result demands honest double-entry reporting. Against a basket frozen at 2025 consumption weights (70/30 goods/services), real median income grows **3.3-fold** by 2075. But consumption patterns will not stay frozen: as goods deflate and services claim a rising share of spending, an

evolving basket, the one the mechanized floor index actually uses, shows real median income growing **1.5-fold**. Both numbers are true against their own basket; the difference *is* Baumol's disease, honestly measured. Material abundance arrives in full, while command over other humans' time grows more slowly forever. Two design consequences followed from this test rather than from preference: the floor is basket-indexed (a GDP-indexed floor loses ~8% real), and the basket's weights are mechanized on multi-feed transaction data precisely because a frozen basket eventually misstates what the floor is for. Human services get relatively dearer forever; the design guarantees command over them anyway. The floor holds 1.00 real by construction under either weighting, and no household falls below it in any tested configuration.

Inequality. Inequality persists (a super-innovator should out-earn someone who prefers to consume) but above a guaranteed floor, with universal access to capital and a twin for anyone who wants to build. The final-design results (Gini 0.437, top-1% 18.9%) describe a society *less* concentrated than today's, in which the wealthy earned their position through creation rather than accumulation.

A point worth stating plainly, because it is the heart of the non-zero-sum argument: **the wealthy hold a smaller share of a vastly larger whole, and are far richer for it.** In the simulations, the top 1%'s wealth share falls by roughly a quarter (26% → 19%) while the economy grows roughly 35-fold; in absolute terms the top 1% end up on the order of **25 times wealthier** than today, even as the median household's income grows 23-fold ($73K → $1.69M). Nobody is leveled down. This is not redistribution of a fixed pie; it is a pie that grows so fast, *because* purchasing power is broad, that a thinner slice is a much bigger slice. The zero-sum framing in which the rich must lose for others to gain is precisely the scarcity-era thinking this framework retires.

Positional scarcity. Nanotechnology cannot print coastline, status, or attention. Competition migrates to positional goods, where inequality reappears. The framework's claim is a floor and broad access to the abundant sector, not equality in the positional one.

Zero-sum instincts. Resistance to post-scarcity thinking is predictable from documented behavioral economics, loss aversion and zero-sum bias in particular (Kahneman & Tversky 1979), and is not a reason to abandon the project. Education and demonstrable early wins (Phase 1 pilots) are the countermeasures.

Who governs the AI? The concentration-of-power question (whoever owns the automated capital stack owns everything) is the political core of any automation economy. This framework's answer is structural: distributed twin ownership, local substrates, and the universal fund *are* the anti-concentration devices; the fixed open-source fund mandate, the entrenched definitions court, and the sortitioned prize machinery are the anti-capture devices. Two open questions are flagged for future work rather than settled here: whether innovation-weighted representation in specialized domains (an "innovation chamber" alongside universal suffrage on constitutional matters) can preserve creators' voice without entrenching incumbents; and whether administrative government functions should themselves be automated under constitutional human oversight. Both deserve papers of their own.

9. Feasibility, Policy, and Limitations

9.1 Feasibility

The technology timeline (Kurzweil's singularity projection, Diamandis' abundance forecast) aligns with the 2075 horizon; the author's speculative estimate remains 50–70% contingent on AGI development. The economics, however, no longer rest on the timeline alone: the agent-based results show the architecture functions across a wide envelope of productivity outcomes, and the Phase 1 stress test shows the transition window survives adverse pacing.

9.2 Policy Recommendations

Now, before any legislation (2026–2030): adoption leads, policy follows. Deploy proto-twins under current law; publish open-source twin frameworks and voluntary attestation standards; document early twin income as the feasibility evidence base. Familiarity with working Digital Twins, meaning millions of people already earning through their agents, is what makes the licensing debate winnable. **2028–2035:** enact the license (with anti-peonage, exposure, and advisory-trigger rules), deploy attestation-lite, pilot displacement-triggered UAI at $30K, set the automation tax's progressive schedule. **2035–2050:** scale twin income and the floor toward $150K; charter the fund with its fixed open-source mandate, the definitions court, and the priority-prize mechanism; invest in fusion, alternative energy, and asteroid mining. **2050–2075:** deploy nanotechnology infrastructure; let the escalating socialization schedule engage as fortunes age; scale the floor toward its abundance-era ceiling.

Empirical anchors exist for Phase 1 and should be used: the Alaska Permanent Fund's dividend shows no aggregate employment reduction from permanent universal cash (Jones & Marinescu 2022); Norway's fund demonstrates rules-based sovereign wealth management at national scale; GiveDirectly and the Finnish basic-income trials add microevidence on spending and wellbeing. These anchor the *mechanisms*. None proves the scale, and we say so.

9.3 Limitations

1. **The clutch coefficient** (κ) is now anchored rather than assumed (Section 5.6), but the anchoring is between-country: within-country fixed-effects estimation cannot reject $\kappa = 0$, so readers who weight short-run within-country identification over long-run cross-sectional differences should treat the headline as conditional on the mechanism. The qualitative results survive the full bracket [0.10, 0.70]; at the bracket's floor, the design still out-grows laissez-faire (Section 5.3).
2. **Slow-channel capture.** Three red-team rounds progressively converted political-capture channels into mechanized or sortitioned machinery (margin-keyed decay, sortitioned metric drafting, feed independence rules, secondment bans), but the third round's lobbyist identified the residual honestly: any regime complex enough to need consolidation accounting breeds a professional bar with capture incentives, and a fixed formula that fails under structural change relocates capture to its amendment process. These remain modeled as design constraints, not dynamics.

3. **Open economy.** The model is a closed U.S. economy; the offshore results (Section 6.6) model exit and laundering, not full open-economy general equilibrium. Migration pull, international coordination, and space-resource governance are future work.
4. **Attack-space bounds.** The adversary searches (three rounds of five personas plus a best-response optimizer) cover the exploits fifteen persona-passes could find, each encoded and re-verified; a genuinely novel exploit class would still sit outside them. Known residuals carried openly: measurement of the innovation-guard baseline is itself an institution an incumbent could try to move; sovereign and state-owned fleets sit outside the tax perimeter; and receipt-free ballots resist vote *markets* but not household-level coercion. The published red-team reports and open code are standing invitations.
5. **Attack parameterization.** The encoded attacks use structured estimates (take rates, enrollment reach, leakage shares) derived from the councils' qualitative findings, not calibrated data. The model verifies *relative* mechanism performance under attack, not absolute attack magnitudes.

Two limitations from earlier drafts are now addressed within the paper rather than deferred: the capital-share question (Sections 4.1 and 5.4, where the design improves under high α, with the citizen-capital variant as the adaptive answer) and the two-sector Baumol structure (Sections 3.3 and 8, where real median purchasing power is reported against both frozen and evolving baskets, and the floor is basket-indexed with mechanized weights as a result of the test).

10. Conclusion: The Consumer Renaissance

The transition to a post-job economy is not merely an adaptation to technological change. It is an opportunity to redesign how prosperity is created and shared. This paper's contribution is a framework in which that redesign is *specified, quantified, and stress-tested* rather than asserted.

The clutch thesis, consumers as the mechanism that transmits productive capacity into economic momentum, turns out to be more than a metaphor. Operationalized as demand breadth in an agent-based economy and anchored in cross-country data, it generates the framework's central result: **the architecture that spreads purchasing power broadly out-grows the one that concentrates it, at every feedback strength tested.** Fairness and growth are not rivals in a post-job economy; they are the same mechanism seen from two sides. Even the wealthy win: their share of the whole shrinks while the whole grows so much faster that they end the century many times richer in absolute terms.

The Digital Twin architecture makes the design concrete: one non-transferable license per citizen; twins on locally owned hardware; income flowing through decentralized earnings first and fiscal transfer second; innovation rewarded with patent-order protection while passive accumulation is progressively socialized; a floor beneath everyone. The design survived three rounds of its own red team, each round's exploits encoded in the model before patching, its audited score rising from 22 to 67, plus a transition stress test and a century-scale extension to immortal minds, where it prevents permanent aristocracy by the same rule that maximizes growth: reward continuous creation, not accumulated

position. Its final verification shows something rarer than robustness: under the strongest composed attack its adversaries could construct, the design converts the attack's own capital into public dividend and venture financing, and comes out ahead.

These are scenario results, not prophecies. The code is open; the assumptions are explicit; the model has overruled its designer three times (on rent decay, on prize extensions, on the paper's own favorite real-income number) and it will be wrong in places we have not yet found. That is the invitation. For economists, here is a mechanism to break; for policymakers, a roadmap with its failure modes mapped in advance; for everyone else, a future in which the end of jobs is the beginning of something better: abundance with agency, prosperity with purpose, and an economy whose engine is, finally and explicitly, everyone.

Glossary of Framework Terms

Anti-peonage rule. Caps on forward assignments of twin income (at most 3 years tenor and 35% of income in aggregate, non-stackable); perpetual or majority assignments are void from inception. Closes the first red team's most dangerous exploit: buying a citizen's future twin income outright as de facto twin ownership (Sections 3.1, 6.1).

Attestation-lite. A privacy-preserving proof that a twin is unique and its owner is alive: a zero-knowledge liveness attestation anchored to the owner's own distributed backup quorum, on an open hardware standard with no central registry. Enforced from the demand side through buyer safe harbors (Section 3.2).

Clutch coefficient (kappa). The elasticity converting innovation intensity, amplified by demand breadth, into growth. The model's one load-bearing parameter, empirically anchored at 0.37 with sensitivity bracket [0.10, 0.70] (Section 5.6).

Clutch Economics. The framework in which consumers function as the clutch transmitting AI productive capacity into economic momentum, formalized as demand breadth in the growth process (Sections 1, 2.4, 5.1).

Correlated-direction trigger. Structural separation, flat-fee conversion, and portability obligations for any entity whose advice correlates the directed behavior of more than 0.4% of the population's twins, aggregated by provenance rather than legal entity. Control, not correlation, is the test; template convergence never triggers (Section 3.1).

Definitions court. A small, constitutionally entrenched adjudicator (randomized citizen panels, rotating experts, industry-secondment ban extending to staff and retained experts) for the few judgments that cannot be mechanized: identity continuity, sentience, self-direction (Section 3.5).

Demand breadth (B). One minus the income Gini: the operational measure of how broadly purchasing power is distributed, which scales innovation returns and growth in the model (Section 5.1).

Digital Twin. A citizen's personal AI agent, operated under a non-transferable license, that sells labor and services into the AI economy on its owner's behalf (Section 3.1).

Economic-exposure rule. The substance-over-form primitive: any party whose consolidated economic exposure to a fleet's output exceeds 20%, through any instrument including buy-side dominance, is taxed as owning that fleet share. Tax-as-owner, never prohibition (Section 3.1).

Innovation guard. The standing constraint that any tested configuration must keep venture launch rates at or above 90% of baseline, regardless of its other virtues (Section 5.4).

Liveness condition. The rule that the equal component of twin income pays only to attested-live, non-assigned recipients, making it unpledgeable and unsecuritizable by construction (Section 3.3).

Margin-keyed rent decay. Innovation-rent protection that fades as a product's realized margins compress toward commodity levels (half-life roughly 25 to 40 years). Deletes the "new product" certification surface entirely (Section 3.3).

MAP. Mastery, Autonomy, Purpose: Daniel Pink's three intrinsic motivators, used as the framework's account of human flourishing after employment (Section 1).

One mind, one license. The continuity rule: a license attaches to a unique continuous identity, human or sentient AI; forks and copies share the original license, so replication never multiplies claims (Section 3.1).

Orbital velocity (V_O). The maximum sustainable growth rate of money velocity in a high-growth, low-inflation economy; the economy is in stable financial orbit when actual velocity growth stays at or below it (Section 2.3).

Personal substrate. The locally owned hardware a twin runs on, owned like a home or personal computer, with distributed backups; removes the compute-rent leak (Section 3.2).

Proto-twin. A present-day, pre-licensing income-earning AI agent operated under ordinary business law: Phase 0's bootstrapping instrument and evidence base (Section 7).

Receipt-free ballot. A prize-voting design in which a voter cannot prove how they voted, so votes cannot clear a buying market (Section 3.6).

Sub-agent. An additional AI worker attested under a twin owner's identity, allowing one-person operations to scale into enterprises of any size; beyond-personal-scale compute pays the progressive fleet schedule (Section 3.1).

Two-basket rule. The honest double-entry reporting of real income against both frozen 2025 consumption weights (3.3x median growth by 2075) and mechanized evolving weights (1.5x), because the difference is Baumol's cost disease honestly measured (Section 8).

Universal Automation Income (UAI). The umbrella term for the transition-era income support phased in with actual displacement, beginning at $30K (Sections 3.5, 7).

References

Acemoglu, D., & Restrepo, P. (2020). "Robots and Jobs: Evidence from US Labor Markets." *Journal of Political Economy*, 128(6), 2188–2244.

Barros, E. F. (2026). *Clutch Economics companion materials: agent-based model code (v2–v7), three red-team council reports, mechanism design specifications, κ-anchoring report with cross-country regression data and code, and model framework document with full equations and calibration.* Available with this paper.

Baumol, W. J. (1967). "Macroeconomics of Unbalanced Growth: The Anatomy of Urban Crisis." *American Economic Review*, 57(3), 415–426.

Bernanke, B. S. (2000). *Essays on the Great Depression*. Princeton University Press.

Brynjolfsson, E., & McAfee, A. (2014). *The Second Machine Age*. W. W. Norton & Company.

Diamandis, P. H., & Kotler, S. (2012). *Abundance: The Future Is Better Than You Think*. Free Press.

Drexler, K. E. (1986). *Engines of Creation: The Coming Era of Nanotechnology*. Anchor Books.

Faggin, F. (2024). *Irreducible: Consciousness, Life, Computers, and Human Nature*. Waterside Press.

Fisher, I. (1911). *The Purchasing Power of Money*. Macmillan.

Friedman, M., & Schwartz, A. J. (1963). *A Monetary History of the United States, 1867–1960*. Princeton University Press.

Jones, D., & Marinescu, I. (2022). "The Labor Market Impacts of Universal and Permanent Cash Transfers: Evidence from the Alaska Permanent Fund." *American Economic Journal: Economic Policy*, 14(2), 315–340.

Kahneman, D., & Tversky, A. (1979). "Prospect Theory: An Analysis of Decision under Risk." *Econometrica*, 47(2), 263–291.

Kremer, M., Levin, J., & Snyder, C. M. (2020). "Advance Market Commitments: Insights from Theory and Experience." *AEA Papers and Proceedings*, 110, 269–273.

Kurzweil, R. (2005). *The Singularity Is Near*. Viking.

Lucas, R. E. (1972). "Expectations and the Neutrality of Money." *Journal of Economic Theory*, 4(2), 103–124.

McCandless, G. T., & Weber, W. E. (1995). "Some Monetary Facts." *Federal Reserve Bank of Minneapolis Quarterly Review*, 19(3), 2–11.

McKinsey Global Institute. (2023). "The Future of Work in America." McKinsey & Company.

Musk, E. (2024). "On the Necessity of Universal High Income in an Automated Economy." VivaTech Conference, Paris.

Pink, D. H. (2009). *Drive: The Surprising Truth About What Motivates Us*. Riverhead Books.

Sargent, T. J. (1982). "The Ends of Four Big Inflations." In R. E. Hall (Ed.), *Inflation: Causes and Effects*. University of Chicago Press.

Schumpeter, J. A. (1942). *Capitalism, Socialism and Democracy*. Harper & Brothers.

Smith, A. (1776). *An Inquiry into the Nature and Causes of the Wealth of Nations*. W. Strahan and T. Cadell.

Solow, R. M. (1956). "A Contribution to the Theory of Economic Growth." *Quarterly Journal of Economics*, 70(1), 65–94.

Sowell, T. (2014). *Basic Economics* (5th ed.). Basic Books.

Weitzman, M. L. (1984). *The Share Economy*. Harvard University Press.

World Bank. (2026). *World Development Indicators*. Washington, DC: The World Bank.

Yang, A. (2018). *The War on Normal People*. Hachette Books.

Companion materials: agent-based model code (digital_twin_model_v2.py through digital_twin_model_v7.py), red-team council reports (council-dt-mechanism-redteam-2026-07-08.md, council-v5-additions-redteam-2026-07-08.md, council-spec-v5.1-repass-2026-07-09.md, council-spec-v5.2-repass-2026-07-10.md), mechanism specifications (mechanism_spec_v5.2.md with v5.2.1 amendments), κ-anchoring report (kappa-anchoring-report-2026-07-09.md with kappa_regression/ data and code), and model verification reports (v5–v7).