Demand Breadth and Innovation-Driven Growth in a Highly Automated Economy: An Agent-Based Analysis

Emanuel F. Barros[1]

Independent Researcher

July 2026

Abstract

If artificial intelligence automates most production, household income becomes a policy variable rather than a labor-market outcome, and a classical question returns with new force: does the *distribution* of purchasing power affect the rate of innovation-driven growth? This paper studies the question in an agent-based model of a highly automated economy in which innovation is endogenous and the return to launching a new product depends on demand breadth, defined as one minus the income Gini coefficient. Two income architectures are compared: a license-based architecture that distributes AI-generated earnings capacity universally while preserving long innovation rents, and a laissez-faire benchmark in which AI earnings accrue to capital owners. Across every tested value of the demand-feedback parameter κ, the broad-distribution architec-

[1]Independent Researcher. Email: efbarros@gmail.com. Replication materials (model code, regression data and code, and adversarial-review reports) accompany this paper; see the Replication Package section. The author received no external funding and declares no conflicts of interest.

ture produces both faster growth and lower inequality than the benchmark, because broad purchasing power enlarges the innovation prize while universal capital access enlarges the pool of potential innovators. The load-bearing parameter κ is disciplined empirically: literature triangulation and a cross-country panel regression center it at $\kappa \approx 0.37$, though within-country fixed-effects estimates cannot reject zero, a limitation stated prominently. The design is further subjected to an unusual robustness exercise: three successive adversarial red-team rounds whose discovered exploits (income-stream purchase, gray AI labor, threshold sharding, monopsony structures, offshore laundering) are encoded as adaptive strategies inside the model before patching. The final architecture is robust to the strongest composed attack found, and the innovation rate never falls below 90 percent of its unregulated baseline in any configuration.

JEL codes: C63, D31, E25, O31, O33, O41.

Keywords: automation; innovation; income distribution; market size; agent-based computational economics; mechanism design.

1 Introduction

Suppose the automation optimists are right about capability and the automation pessimists are right about labor: within some decades, artificial intelligence and robotics perform most tasks that wage labor performs today, and the wage channel through which production has historically financed consumption narrows toward irrelevance. Nearly all analysis of this scenario concerns its distributional consequences (Korinek and Stiglitz, 2019; Moll et al., 2022). This paper asks a different question, about growth: *in an economy where most income is machine-generated and therefore institutionally assignable, does the assignment itself, the breadth or concentration of purchasing power, affect the rate of innovation-driven growth?*

The question is old. Murphy et al. (1989) showed that industrialization requires a middle class large enough to buy manufactured goods; Foellmi and Zweimüller (2006) and Zweimüller (2000) formalized how income distribution shapes demand-induced innovation when new goods diffuse down Engel curves; Matsuyama (2002) modeled the virtuous circle between mass consumption and productivity growth. The empirical market-size literature gives the mechanism causal footing: pharmaceutical innovation responds strongly to expected market size (Acemoglu and Linn, 2004; Dubois et al., 2015; Costinot et al., 2019), product entry tracks the purchasing-power growth of the customer base (Jaravel, 2019), and invention has followed demand at least since Schmookler (1966). What automation adds is leverage. In a wage economy, the distribution of purchasing power is largely determined by the labor market, and policy adjusts it at the margins. In a highly automated economy, distribution becomes a design choice, and if the demand-breadth channel is real, that choice feeds back into the growth rate itself.

We study this feedback in an agent-based model, in the computational-economics tradition of Tesfatsion (2006) and Farmer and Foley (2009), with three features that jointly distinguish it from existing automation models. First, innovation is endogenous: heterogeneous agents launch ventures when the expected present value of innovation rents justifies it, and the aggregate innovation rate drives productivity growth. Second, the returns to innovation depend on demand breadth $B_t = 1 - \mathrm{Gini}(y_t)$ through a single feedback parameter κ, which is the model's formal statement of the demand-breadth hypothesis. Third, the policy experiment is institutional rather than fiscal: we compare a *license architecture*, in which each adult holds one non-transferable license to an AI agent earning market income on the holder's behalf, with a laissez-faire benchmark in which AI earnings accrue to the owners of AI capital. The license architecture is a stylization of one concrete proposal (Barros, 2026), but the results speak to any institution that distributes AI-era earnings capacity broadly, universal capital endowments and sovereign-wealth dividends among them.

Three findings emerge. The first is the paper's central result: **the broad-distribution architecture out-grows laissez-faire at every tested value of κ**, while producing income Gini coefficients of 0.34 to 0.42 against 0.70 to 0.79 under the benchmark (Section 4). The mechanism is not redistribution of a fixed surplus. Two channels do the work: demand breadth enlarges the prize per venture, and universal earnings capacity enlarges the set of agents who can attempt ventures at all. Under the benchmark, both channels atrophy as income concentrates. Growth and equality are complements, not substitutes, in this environment; the model gives conditions under which the classic equity-efficiency trade-off inverts.

Second, because a sensitivity sweep shows κ is the only materially load-bearing free parameter (Section 4.4), we discipline it empirically rather than assume it (Section 5). Mapping the parameter onto adjacent estimated quantities in the market-size, marginal-propensity-to-consume, and inequality-growth literatures brackets $\kappa \in [0.1, 1.0]$ with a central band of 0.3 to 0.5. A cross-country panel regression on World Bank data in the model's own functional form yields a between-country estimate of $\kappa \approx 0.37$ ($p = 0.004$). We state the identification problem as prominently as the estimate: within-country fixed-effects estimation cannot reject $\kappa = 0$, the familiar pattern in which between and within estimators disagree in the inequality-growth panel literature (Forbes, 2000; Barro, 2000; Halter et al., 2014).

The paper's growth-advantage claim is therefore conditional on cross-sectional identification; readers who insist on within-country identification should read the results as distributional dominance at growth parity, which is how the model behaves as $\kappa \to 0$.

Third, the paper contributes a robustness exercise we believe is rare in the mechanism-design literature: **strategic-manipulation robustness by adversarial iteration** (Section 6). Any institution that assigns income invites gaming, and paper designs are usually tested against a referee's imagination rather than a systematic search. We conducted three successive structured red-team rounds against the license architecture, five adversarial personas per round; each round's discovered exploits (purchase of future income streams, unlicensed gray AI labor, wealth sharding across tax thresholds, monopsony reconstruction through output purchases, offshore laundering of AI output) were encoded as adaptive strategies inside the model, run against both the unpatched and patched designs, and only then patched. The unpatched design degenerates into re-concentrated oligopoly, as the red team predicted. The final architecture holds: under the strongest composed attack discovered, twin-income independence remains at 83 percent, inequality returns approximately to baseline, and the attack is converted into public revenue. Throughout, every configuration is required to satisfy an *innovation guard*, a constraint-based welfare criterion under which any policy that pushes the venture launch rate below 90 percent of its unregulated baseline fails regardless of its distributional virtues.

The paper relates to four literatures. The automation literature establishes displacement and factor-share effects (Acemoglu and Restrepo, 2018, 2020; Autor, 2015; Frey and Osborne, 2017) and studies distributional dynamics under automation (Moll et al., 2022; Korinek and Stiglitz, 2019); we take its displacement scenario as given and endogenize the innovation response to the induced income distribution. The demand-composition literature cited above supplies the theoretical mechanism; our contribution is quantitative, embedding it in a full economy with endogenous entry and asking whether it survives calibrated feedback and adversarial pressure. The empirical inequality-growth literature supplies both our anchoring strategy and our humility: reduced-form estimates disagree by estimator (Cingano, 2014; Ostry et al., 2014; Berg et al., 2018; Brueckner and Lederman, 2018; Banerjee and Duflo, 2003; Neves et al., 2016), and we inherit that disagreement honestly. Finally, the heterogeneous-MPC literature (Jappelli and Pistaferri, 2014; Carroll et al., 2017; Kaplan and Violante, 2014) grounds the assumption that broad income translates into broad spending. On policy instruments, the license architecture's dividend component connects to the one large-scale universal cash transfer with credible causal evidence, the Alaska Permanent Fund dividend, which shows no aggregate employment reduction (Jones and Marinescu, 2022), and its prize component to advance market commitments (Kremer et al., 2020).

Two disclaimers frame everything that follows. All quantitative results are *scenario analyses* conditional on an assumed automation trajectory, not forecasts; the model's arithmetic is computationally verified, but its premises are assumptions. And the paper's normative content is limited to a single constraint we impose on ourselves: whatever else an income architecture does, it must not materially suppress innovation. Speculative extensions of the framework, including long-horizon abundance scenarios, are confined to Appendix A and to the companion working paper (Barros, 2026).

Section 2 presents the model, Section 3 the calibration, Section 4 the main results, Section 5 the empirical anchoring of κ, Section 6 the strategic-manipulation robustness exercise, Section 7 limitations, and Section 8 concludes.

2 The Model

The model is an agent-based simulation of a closed economy over 2025 to 2075, populated by N households scaled to the United States adult population. It is deliberately parsimonious: rather than a full general-equilibrium treatment, it isolates the interaction of four blocks (automation, income assignment, endogenous innovation, and demand feedback) and subjects that interaction to Monte Carlo and adversarial analysis. Full pseudocode and the complete parameter list ship with the replication package.

2.1 Environment and automation

Time is annual. An automation index $a(t)$ rises along a logistic path from near zero in 2025 toward one by mid-century, shifting production from wage labor to AI capital and calibrating the disappearance of wage income. Aggregate output Y_t grows at an endogenous rate g_t specified in Section 2.4. A guaranteed income floor ($40,000 per household in 2025 terms, indexed as described below) is financed from the architecture's fund; the share of households below the floor is tracked as a failure metric and equals zero in all baseline configurations.

2.2 Households

Each household i carries three state variables: wealth w_{it}, initialized from a calibration of the 2025 U.S. wealth distribution (Section 3); innovation ability θ_i, drawn from a heavy-tailed distribution so that a small right tail of agents supplies most high-value ventures; and a portfolio of active innovation rents from past successful launches. Households receive income through the channels their architecture provides (Section 2.5), consume with a marginal propensity declining in wealth, consistent with the evidence in Carroll et al. (2017) and Jappelli and Pistaferri (2014), and accumulate the remainder as wealth earning a passive return r.

2.3 Endogenous innovation

Innovation is the model's engine and its central margin of interest. In each period, household i launches a venture with probability

$$p_{it}^{\text{launch}} = \bar{p} \cdot \theta_i \cdot \frac{B_t/(r+\bar{\delta})}{\text{NPV}_{\text{ref}}} \cdot \mathbb{1}[\text{capital access}_{it}]\,, \tag{1}$$

where $\bar{\delta}$ is the decay rate of innovation rents (so $1/(r + \bar{\delta})$ is the present-value factor of a rent stream), NPV_{ref} is a normalizing reference value, and B_t is demand breadth, defined below, which scales the expected prize: a broad customer base raises the value of a new product, a concentrated one lowers it. The capital-access indicator is the institutional margin: under the license architecture every household can finance a launch (from twin income, the dividend, or safe-harbored outside capital as in Section 6); under laissez-faire, access is wealth-gated. A successful launch endows the household with a rent stream that decays at rate $\bar{\delta}$, corresponding to a protection half-life $h = \ln 2/\bar{\delta}$; rents that have decayed to commodity margins flow into the economy's commoditized profit pool.

Equation (1) embeds the two mechanisms that drive all of the paper's comparative results: the *prize channel* (B_t raises expected returns per venture) and the *participation channel* (capital access determines how much of the ability distribution can respond to those returns). Both are standard in the demand-induced innovation literature (Foellmi and Zweimüller, 2006; Murphy et al., 1989); the model's contribution is to let them interact dynamically with the income distribution the architecture generates.

2.4 Demand breadth and growth

Demand breadth is measured directly on the model's income distribution:

$$B_t \;=\; 1 - \text{Gini}(y_{1t}, \ldots, y_{Nt}). \tag{2}$$

Aggregate growth is

$$g_t \;=\; g_0 \;+\; \kappa \cdot \frac{\sum_j R_{jt}}{Y_t} \cdot (0.5 + B_t) \;+\; \gamma\, a(t), \tag{3}$$

with baseline drift $g_0 = 0.005$, automation productivity term $\gamma = 0.015$, and $\sum_j R_{jt}/Y_t$ the economy's active innovation-rent intensity, which ranges over roughly 0.05 to 0.15 in simulation. The middle term is the demand-breadth hypothesis in one line: innovation-driven growth scales with the interaction of innovation intensity and the breadth of purchasing power, governed by κ. At $\kappa = 0$ the distribution of income has no growth consequence and the model collapses toward a standard automation-growth scenario; everything the paper claims about growth rides on $\kappa > 0$, which is why Section 5 is devoted to estimating it. The functional form, innovation intensity interacted with $(0.5 + B_t)$, keeps the demand channel bounded (halving to doubling the innovation term as B_t spans its range) and is the same form taken to the data.

2.5 The two architectures

License architecture (proposed). Each adult holds one non-transferable license to a personal AI agent (a "twin") that earns market income on the holder's behalf. Licensed twins collectively receive a majority share of the AI-labor compensation pool, sustained by a progressive compute tax on unlicensed datacenter-scale fleets. Twin income has an equal component (half the licensed pool, split per license) and a performance component (half, allocated by direction skill, proxied in-model by θ_i and past success). Three further flows complete the architecture: a universal dividend paid by a passively managed fund that receives commoditized profits and a progressive share of large passive returns (50 percent of passive returns above $10 million, escalating); innovation rents with the long half-life of Section 2.3, never subject to socialization while active; and the floor, topped up from the fund when a household's other channels fall short. The design principle, examined at length in the companion paper, is to socialize only commoditized and passive returns while leaving creation rents untouched.

Laissez-faire benchmark. AI-labor compensation accrues to AI capital owners in proportion to wealth; there is no license, dividend, or floor; venture finance is wealth-gated. The benchmark is not a straw man: it receives the same automation productivity term, the same innovation technology, and the same rent protection, and it generates substantial growth. It differs only in the assignment of machine earnings.

Both architectures are run under identical draws of ability, initial wealth, and shocks. All reported statistics are Monte Carlo medians across seeds.

2.6 The innovation guard

Because any income architecture can purchase equality at the price of dynamism, we impose a constraint-based welfare criterion on every experiment: the *innovation guard*. A configuration fails if its venture launch rate falls below 90 percent of the launch rate in the unregulated baseline economy. The guard converts "does redistribution kill innovation?" from a rhetorical objection into a reported, falsifiable statistic; guard status is reported for every experiment in the paper, including every adversarial configuration in Section 6.

3 Calibration and Computational Design

Table 1 summarizes the calibration. The 2025 baseline approximates the United States: GDP $25 trillion, 250 million adults, initial wealth calibrated to the U.S. wealth distribution's documented concentration (top-1 percent share of 26 percent), and a floor of $40,000. The simulation uses $N = 6{,}000$ agents with populations and dollar aggregates scaled to national totals, six Monte Carlo seeds per configuration, and an annual step over 2025 to 2075. Innovation-rent protection uses the half-life range selected by the model itself in Section 4.2 (25 to 40 years, patent-order durations). The demand-feedback parameter is set to the empirically anchored $\kappa = 0.37$ (Section 5) with the full bracket $[0.10, 0.70]$ swept in robustness runs; the headline comparative result of Section 4.1 is reported across the sweep precisely so that no single κ carries it.

Three calibration choices deserve comment. First, the floor is indexed to a consumption basket with mechanized, evolving weights rather than to GDP: a two-sector extension of the model with Baumol dynamics (Baumol, 1967), in which goods deflate while human-provided

Table 1: Calibration summary.

Block	Parameter	Value (source)
Baseline	GDP 2025	\$25T (U.S., approx.)
	Adults	250M
	Top-1% wealth share, 2025	26% (U.S. distribution)
	Floor	\$40K, basket-indexed
Automation	$a(t)$ path	logistic, near-complete by mid-century
	Productivity term γ	0.015
Innovation	Ability θ_i	heavy-tailed (right-skewed)
	Rent half-life h	25–40 yr (model-selected, §4.2)
	Baseline drift g_0	0.005
Demand feedback	κ	0.37, bracket $[0.10, 0.70]$ (§5)
Architecture	Twin pool share	≈60% of AI-labor compensation
	Equal/performance split	50/50
	Socialization	50% of passive returns >\$10M, escalating
	Fund payout	95% of income as dividend, 5% buffer
Computation	Agents N	6,000 (scaled)
	Seeds	6; medians reported
	Horizon	2025–2075, annual

Full parameter list, distributional calibration targets, and pseudocode ship with the replication package (Section 8).

services remain wage-indexed, shows that a GDP-indexed floor loses roughly 8 percent of real purchasing power by 2075 and a frozen-weight basket drifts further; the indexation rule is a result, not an assumption, and Section 6.4 reports the honest double-entry consequence for real income claims. Second, wage income disappears along $a(t)$ by construction; readers who reject near-complete automation should read all results as conditional on that premise. Third, the model is closed: offshore exit appears only as an adversarial strategy (Section 6), not as general-equilibrium trade.

4 Results

4.1 Demand breadth versus concentration

Table 2 and Figure 1 present the central experiment: the license architecture against laissez-faire across the demand-feedback sweep, identical technology and shocks, Monte Carlo me-

Table 2: License architecture versus laissez-faire across demand-feedback strengths.

	Real GDP 2075 ($T)		Income Gini 2075		Below floor	
κ	License	Laissez-faire	License	Laissez-faire	License	LF
0.10	144	117	0.42	0.70	0%	n/a
0.20	306	193	0.39	0.73	0%	n/a
0.35	786	351	0.36	0.77	0%	n/a
0.50	1,881	527	0.34	0.79	0%	n/a

Monte Carlo medians. Identical ability draws, initial wealth, and shocks across architectures. "Below floor" is the share of households under the guaranteed floor; the benchmark has no floor. Absolute GDP levels are conditional on the automation scenario and should be read as scenario outputs; the comparative statement (license > laissez-faire at every κ) is the result.

dians.

The license architecture out-grows the benchmark at every feedback strength, with the margin widening in κ: 23 percent higher GDP at $\kappa = 0.10$, 2.2 times at $\kappa = 0.35$, 3.6 times at $\kappa = 0.50$. Decomposing the mechanism by switching channels off separately, both margins matter. With the prize channel disabled (returns independent of B_t), the architectures converge on growth and differ only in distribution. With the participation channel disabled (universal capital access removed from the license architecture), roughly half the growth gap closes: concentrating income does most of its growth damage by shrinking the set of agents who can act on innovation incentives, the same scarcity that motivated Murphy et al. (1989) in the industrialization context. Inequality dynamics are equally stark: laissez-faire concentrates (Gini 0.70 to 0.79 by 2075) because AI earnings compound on wealth, while the license architecture holds Gini near 0.4, with zero households below the floor at every κ.

Two remarks bound the claim. At $\kappa = 0$, untabulated, the architectures tie on growth and the result reduces to distributional dominance; the growth advantage is entirely a demand-breadth phenomenon and stands or falls with Section 5. And absolute levels (a 6- to 75-fold GDP expansion) inherit the automation premise wholesale; we emphasize ratios between architectures, which difference out the shared technology path.

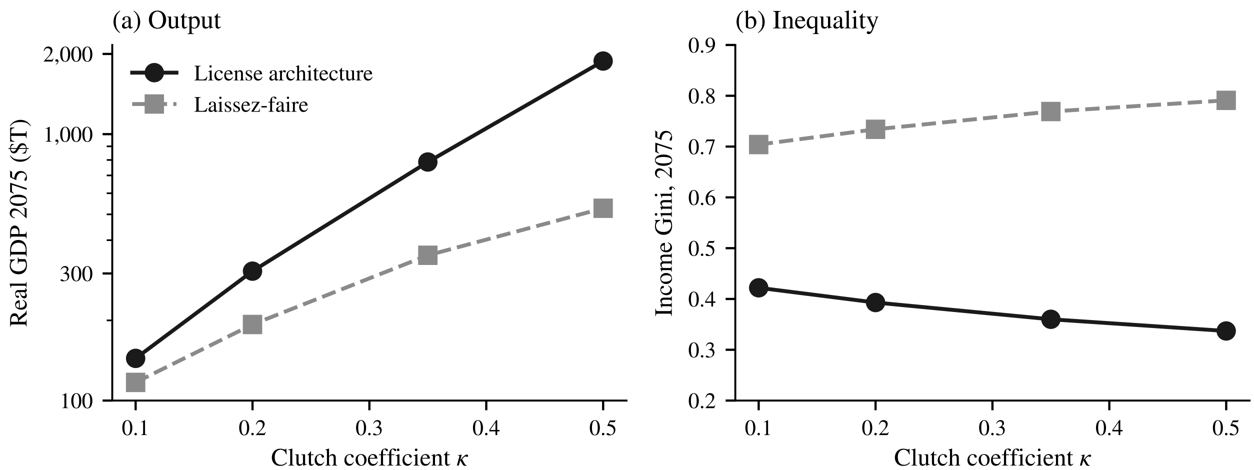

Figure 1: Output and inequality under the two architectures across the demand-feedback sweep. The license architecture (solid) dominates laissez-faire (dashed) on output at every κ while holding income Gini near 0.4 against 0.7 to 0.8 under the benchmark. Monte Carlo medians; log scale in panel (a).

4.2 Rent duration: the model against its designer

How long should innovation rents last? The question matters because rent decay $\bar{\delta}$ enters the launch decision (1) directly. Sweeping the protection half-life from 2 to 40 years (Figure 2) shows GDP rising steeply with duration (from \$69T to \$870T at 2075) while income Gini rises only gently (0.20 to 0.37): long rents power the launch incentive, while the distributional work is done elsewhere, by the license and socialization instruments. The design consequently adopts patent-order durations (25 to 40 year half-life), *overruling the designer's prior*, which had favored aggressive decay; we flag this because a model that never contradicts its author is not doing much work. Distributional instruments and innovation incentives are separable in this architecture: a result consistent with the Schumpeterian principle that growth needs prospective rents (Schumpeter, 1942; Aghion and Howitt, 1992), and one that removes the usual motivation for taxing active innovation returns.

4.3 The capital-share stress test

The automation literature broadly expects capital's share of income to rise as automation proceeds (Acemoglu and Restrepo, 2018; Moll et al., 2022), which would seem to threaten

Figure 2: Rent-protection half-life sweep. Growth rises steeply with rent duration while inequality rises only gently; the license and socialization instruments, not rent decay, carry the distributional load. Monte Carlo medians.

an architecture whose primary channel is quasi-labor income earned by licensed agents. We therefore re-ran the economy with capital's share ramping to 60 percent of output as automation completes, in two variants: the design unchanged, and an adaptive variant adding an equal per-license claim on commoditized capital income.

Table 3 and Figure 3 show the objection inverting. Under the high capital share, the swollen profit pool flows into commoditized inflows, and the universal dividend automatically compensates for shrinking twin labor income: the fund is a built-in stabilizer against exactly the factor-share shift the literature predicts. GDP and median income *rise* relative to the labor-share world, the launch rate rises with them (broader income means bigger prizes), and the guard passes throughout. The honest reading is compositional: under high capital shares the dividend, not the twin, becomes the dominant channel, so the architecture's decentralized-earnings philosophy survives only in the citizen-capital variant, which also performs best on every metric. The factor-share dispute changes which channel dominates, not whether the design works.

Table 3: Capital-share stress test (2075 outcomes, $\kappa = 0.35$).

	Baseline (labor-share world)	High capital share, design unchanged	High capital share + citizen capital share
Real GDP ($T)	635	828	983
Median income ($K)	1,250	1,731	2,117
of which twin ($K)	673	240	280
of which dividend ($K)	134	970	584
of which citizen capital ($K)	–	–	760
Income Gini	0.42	0.31	0.28
Top-1% wealth share	17%	12%	12%
Below floor	0%	0%	0%
Innovation guard	PASS	PASS	PASS

Monte Carlo medians. "High capital share": capital income ramps to 60% of output as automation completes. Guard threshold: launch rate $\geq$ 90% of unregulated baseline.

4.4 Sensitivity: one load-bearing parameter

A tornado sweep varies every assumed parameter by $\pm$50 percent around the baseline (Figure 4). Only κ materially moves output: GDP 2075 spans $213T to $1,517T across its range, roughly three times the swing of the next parameter. The licensed pool share is the dominant *distributional* lever (Gini 0.55 to 0.34 across its range), consistent with Section 4.2's finding that licensing rather than rent decay does the fairness work. No parameter in the sweep drives the launch rate below the innovation guard. This concentration of sensitivity in a single parameter is what motivates the next section: the model's one load-bearing free parameter deserved an estimate, not an assumption.

5 Empirically Anchoring the Demand-Feedback Parameter

This section disciplines κ by two independent routes: triangulation from adjacent estimated literatures, and a direct cross-country regression in the model's functional form. We state the conclusion first, including its weakness: the two routes agree on a central value near 0.37, but identification is cross-sectional. Within-country fixed-effects estimation cannot reject

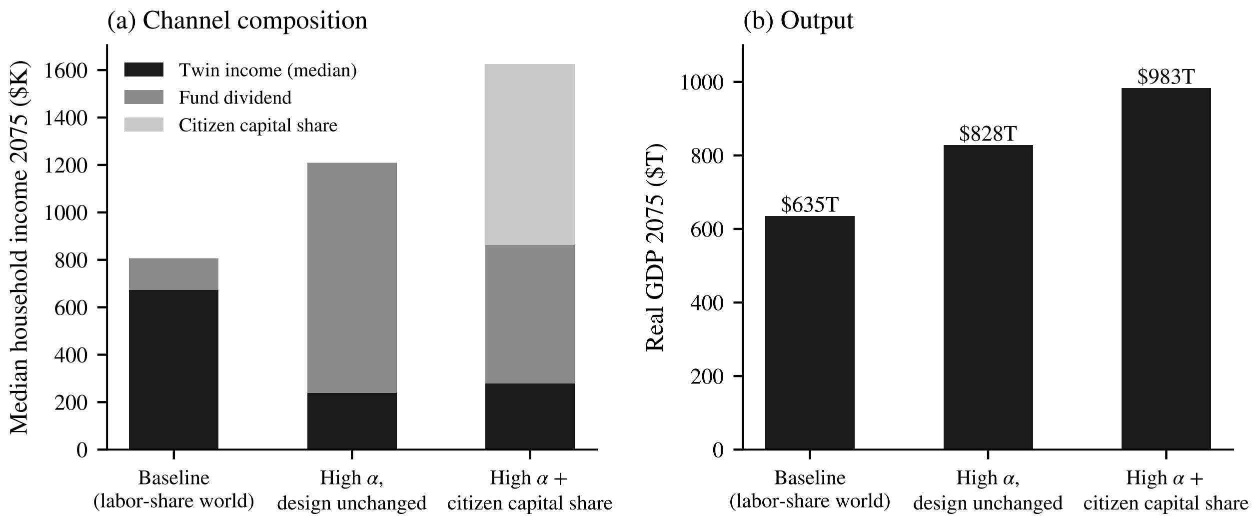

Figure 3: Capital-share stress test. (a) Median income by channel: the dividend absorbs the factor-share shift; the citizen-capital variant restores a direct ownership channel. (b) Output rises under the high capital share. Monte Carlo medians.

$\kappa = 0$, and readers should weight the paper's growth claims accordingly.

5.1 Literature triangulation

The parameter κ converts a change in demand breadth, interacted with innovation intensity, into a growth-rate change. Under the model's operating range (innovation intensity $R/Y \approx 0.10$), moving Gini from the laissez-faire outcome (0.77) to the license outcome (0.36) raises B by 0.41 and growth by 0.041κ per year, so one Gini point is worth roughly 0.1κ percentage points of annual growth. This conversion lets three literatures speak to κ:

Inequality-growth reduced forms. Cross-country estimates of the growth cost of inequality (Cingano, 2014; Ostry et al., 2014; Berg et al., 2018; Brueckner and Lederman, 2018) imply $\kappa \approx 1$ to 2 under the conversion, but these are upper bounds for our purpose: reduced forms bundle every channel from human capital to political instability, while κ captures only demand-breadth amplification of innovation returns. The meta-analytic central estimate (Neves et al., 2016) implies roughly 0.11 to 0.15, and the panel literature's estimator disagreements (Forbes, 2000; Barro, 2000; Halter et al., 2014; Banerjee and Duflo, 2003) caution against leaning on any single reduced-form number.

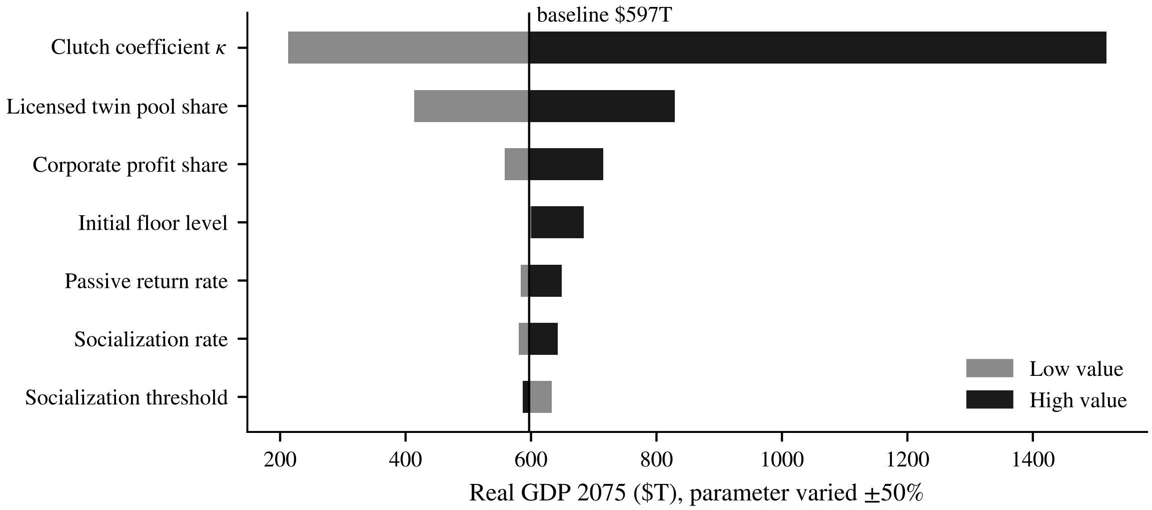

Figure 4: Tornado sensitivity of 2075 GDP to ±50 percent parameter variation. The demand-feedback parameter κ dominates; all other parameters are second-order. No variation breaches the innovation guard.

Market size and innovation. The mechanism-specific literature is closest to κ's definition. Pharmaceutical innovation responds to market size with entry elasticities near or above unity in Acemoglu and Linn (2004) and Costinot et al. (2019), and around 0.23 in Dubois et al. (2015); retail product entry tracks customer purchasing-power growth roughly proportionally in Jaravel (2019). Mapped through the model's rent-intensity range, these imply κ between roughly 0.03 and 0.9, with the causal estimates favoring the middle of that range.

MPC heterogeneity. Marginal propensities to consume of roughly 0.6 for low-wealth households against 0.2 to 0.3 for the affluent (Jappelli and Pistaferri, 2014; Carroll et al., 2017; Kaplan and Violante, 2014) imply that redistributing a dollar raises aggregate spending, a level effect worth roughly $\kappa \approx 0.12$ to 0.17 during transitions.

Jointly: $\kappa \in [0.1, 1.0]$, central band 0.3 to 0.5.

5.2 Cross-country regression

We estimate the model's growth equation directly on World Bank World Development Indicators data (World Bank, 2021): 217 economies, analysis window 1990 to 2019, Gini available for 1,643 country-years across 167 countries. Demand breadth is $B = 1 - \text{Gini}/100$, lagged

Table 4: Growth and demand breadth: panel and cross-section estimates.

Specification	$\hat{\beta}_B$	s.e.	p	N	Implied κ
Within-country (fixed effects, 5-year panels)					
Growth, FE	−0.052	0.048	0.28	504 (141 c.)	n.s.
log patents per million, FE	−0.780	2.103	0.71	390	n.s.
R&D/GDP, FE	−0.709	0.851	0.41	365	n.s.
Growth, FE, OECD	−0.071	0.071	0.32	159	n.s.
Growth, FE, trimmed	−0.113	0.055	0.04	401	negative
Between-country (long-run cross-section, 1995–2019 average growth)					
Cross-section	+0.037	0.013	0.004	146	**0.37**
Trimmed GDP deciles	+0.044	0.013	0.001	116	0.44
OECD only	+0.060	0.027	0.025	36	0.60

Dependent variable: average annual real GDP growth (decimal) except where noted. $B = 1 - \text{Gini}/100$, lagged one 5-year period. Controls: log initial GDP per capita, investment share, secondary enrollment. FE specifications include country and period effects, SEs clustered by country; cross-section SEs are HC1. Implied $\kappa = \hat{\beta}_B/(R/Y)$ at the model's central rent intensity $R/Y = 0.10$; at $R/Y \in \{0.15, 0.05\}$ the cross-section point estimate implies 0.25 and 0.74. Data: WDI, 2021 vintage (World Bank, 2021).

one five-year period; controls are log initial GDP per capita, the investment share, and secondary enrollment; fixed-effects standard errors are clustered by country and cross-section errors are heteroskedasticity-robust. Table 4 reports the specifications.

The pattern is the classic between-versus-within disagreement of the inequality-growth literature. The between-country estimator is robustly positive: ten Gini points of additional breadth associate with roughly 0.37 percentage points of annual growth, implying $\kappa \approx 0.37$ at the model's central rent intensity, surviving trimming and the OECD subsample (Figure 5). The within-country estimators are null across growth, patents, and R&D outcomes, with one marginally negative trimmed cell that we report and decline to lean on in either direction. Both facts have textbook explanations that we cannot rule out: the between estimate is confounded by institutions and state capacity in the expected direction, making it a soft upper anchor rather than a causal estimate, and the within estimate is attenuated because Gini moves slowly within countries and WDI Gini mixes survey concepts. What the data discipline is magnitude, $|\kappa| \lesssim 1$, and the long-run cross-sectional location near 0.37; they do not deliver causal identification, and zero remains inside the within-country credible set.

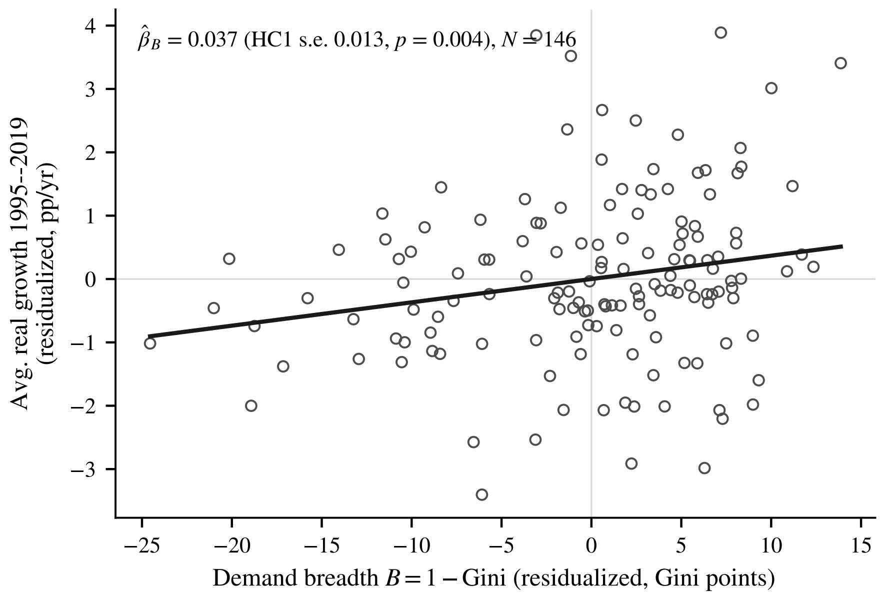

Figure 5: Added-variable plot for the long-run cross-section: average 1995–2019 growth against lagged demand breadth B, both residualized on log initial GDP per capita, investment, and enrollment. The slope equals the cross-section estimate in Table 4.

5.3 Adopted calibration and what rides on it

The two routes converge: literature central band 0.3 to 0.5, regression between-estimate 0.37. We adopt $\kappa = 0.37$ with working bracket $[0.10, 0.70]$, and we re-ran the architecture across the bracket: GDP 2075 ranges from \$129T ($\kappa = 0.10$) through \$740T (0.37) to \$3,175T (0.70), with the guard passing, the floor binding for no one, and Gini between 0.39 and 0.45 throughout. (These bracket runs use the pre-amendment architecture; the final amended design of Section 6 shifts the central level to \$875T at $\kappa = 0.37$, Appendix A.2, without changing any qualitative statement.) Because Table 2 already shows the comparative result at every κ, the paper's central claim needs only $\kappa > 0$; the anchored value determines magnitudes, not signs. Four caveats travel with the anchor: the literature identifies off 5 to 15 Gini-point variation while the model's counterfactual swings 41 points, a 3- to 8-fold out-of-sample extrapolation; linear scaling likely overstates at that distance (Banerjee and Duflo, 2003); all evidence comes from wage-labor economies, and the mapping from measured income Gini to effective demand may differ structurally in a highly automated economy; and mapping pharmaceutical market-size elasticities across all sectors assumes a generality the data do not establish.

6 Robustness to Strategic Manipulation

Mechanism proposals are usually stress-tested against referees' intuitions. We subjected the license architecture to a structured adversarial process closer in spirit to security auditing: three successive red-team rounds, each staffed by five adversarial personas (a regulatory-arbitrage attorney, a platform monopolist, a financial engineer, a political-capture lobbyist, and a shadow economist), each round attacking the design as then specified. Crucially, every discovered exploit was *encoded as an adaptive strategy inside the simulation* and run against patched and unpatched variants before any patch was accepted, so the claim "the patch holds" is a model output, not an assertion. The full council reports, exploit taxonomies, and per-round audit verdicts ship with the replication package; this section reports the quantitative core. We believe the exercise, publishing the attack surface and the failure modes alongside the design, is itself a contribution: a design whose red-team history is public invites attackers to find exploits its adversaries missed, rather than to gesture at the possibility of gaming.

6.1 Round one: buying the income stream

Four of five personas independently converged on the same exploit: the license is non-transferable, but its *output* is transferable property, so purchasing a citizen's future twin income forever, bundled with an "advisory" contract directing the twin, reconstitutes fleet ownership through contract law alone. Companion attacks: unlicensed gray AI labor undercutting the licensed pool, and sharding wealth across the socialization threshold. Encoded and simulated, the unpatched design degenerates exactly as predicted: advisory capture takes 36 percent of twin income, the licensed share of the AI-labor pool collapses to 16 percent, effective socialization falls to a fifth of statutory, and the system lands at Gini 0.73

Table 5: Round-one attacks, encoded and simulated: patched versus unpatched design.

Metric (2075 unless noted)	Target	Patched	Unpatched
Advisory share of twin income (2055)	<10%	5.0%	36.1%
Licensed twins' share of AI-labor pool	≈60%	53.6%	16.4%
Effective socialization (share of statutory)	high	85%	20%
Income Gini	low	0.41	0.73
Top-1% wealth share	bounded	17%	45%
Median income		$1.31M	$283K
Real GDP		$631T	$274T

Monte Carlo medians at $\kappa = 0.35$. Attack strategies: advisory/forward-purchase capture, gray AI labor, threshold sharding, run simultaneously.

with a top-1 percent wealth share of 45 percent, twice the 2025 calibration, at roughly half the patched economy's GDP (Table 5).

The patch set is contractual and structural rather than discretionary: forward assignments of twin income are capped (three-year tenor, 35 percent of income, non-stackable) with larger assignments void; advisory services that correlate the directed behavior of more than 0.4 percent of twins face structural separation and mandatory portability, aggregated by provenance so subsidiary relabeling does not evade the trigger; attestation of licensed uniqueness is demand-side (buyers of attested AI labor receive tax safe harbor) with no central registry; and the socialization threshold applies with beneficial-owner look-through. Under attack, the patched design holds all four survival metrics near their targets, at more than twice the unpatched GDP.

A best-response search over attack intensities within the patched rules (advisory take at its cap, aggressive contract offers, gray-labor investment at attestation limits, maximal sharding within look-through) found graceful degradation as the worst case: top-1 percent wealth 23.3 percent against a 19.5 percent no-attack baseline, Gini 0.49 against 0.44, median income above $1.1M, no one below the floor, guard passing.

6.2 Rounds two and three: attacking the patches

Round two attacked the mechanisms added after round one and found two critical exploits: revenue-contingent financing that rebuilt corporate fleets one contractual layer down, and a prize-duration bonus (part of an earlier design) that recreated a captureable certification surface. The second finding illustrates the model-as-arbiter method: simulated across cash-only, constrained-extension, and unconstrained-extension prize variants, duration bonuses bought *zero* additional innovation over cash prizes while top-1 percent concentration rose monotonically with extension permissiveness, so the mechanism was deleted on cost-benefit grounds. Round two's patches introduced an economic-exposure rule: any party whose consolidated economic exposure to a licensed fleet's output exceeds 20 percent, through any instrument (equity, revenue strips, output-indexed derivatives, contingent leases, or dominant output purchase), is taxed as owning that fleet share, with fixed-fee leases, sub-20-percent equity, and grants as safe harbors. Notably, the model also arbitrated *against* a stricter alternative: banning outside capital entirely would protect independence but locks capital-poor, high-ability agents out of scale, and the exposure-taxed financing channel raises GDP and median income above the no-financing baseline. Financing is an innovation input; the rule taxes control, not capital.

Round three then attacked the exposure rule itself, discovering five structures: manufactured spot monopsony (dominant output purchase without contracts), an agency marketplace holding exposure just under the threshold at scale, matched-book derivative netting hiding net-long exposure, syndicated rate-card stacking, and, most damaging, an offshore spot laundry re-importing unattested AI output. Encoded (Table 6, Figure 6), the offshore laundry proved the worst single channel discovered in all three rounds: 35 percent of the twin pool drains offshore, median income falls to $988K, and it is the only exploit in three rounds to break the innovation guard. The composed pipeline (marketplace plus offshore plus nominee diversification) is super-additive, cutting independent twin income to 45 percent.

The final amendments close the round-three surface: export withholding at the fleet-tax

Table 6: Round-three attacks against the pre-amendment design, and the amended design against the composed attack.

Configuration	GDP ($T)	Median ($K)	Gini	Indep. share	Offshore share	Guard
No attack	693	1,399	0.436	100%	0%	PASS
Spot monopsony	638	1,171	0.469	49%	0%	PASS
Agency marketplace (19.9%)	667	1,269	0.449	100%[a]	0%	PASS
Matched-book netting	722	1,375	0.470	58%	0%	PASS
Rate-card stacking	769	1,353	0.497	70%	0%	PASS
Offshore spot laundry	603	988	0.520	49%	35%	**FAIL**
Composed pipeline	540	885	0.513	45%	40%	PASS[b]
Composed pipeline vs. amended design	767	1,510	0.436	83%	6%	PASS

Monte Carlo medians, $\kappa = 0.37$. "Indep. share": share of twin income earned free of attacker extraction. [a]The marketplace's uniform sub-threshold skim is invisible to the enrollment metric by construction; its damage appears in median income and concentration. [b]At the guard line (0.0426 vs. threshold 0.0421). Amendments: export withholding at non-attested exits (creditable on onshore monetization), aggregate exposure cap across instruments, provenance aggregation of rate-card syndicates.

rate on output exiting to non-attested counterparties (creditable on demonstrated onshore monetization, so no domestic launch is touched), an aggregate exposure cap across instruments, and provenance aggregation of syndicated positions. Against the composed attack, the amended design restores independence to 83 percent, cuts offshore leakage to 6 percent, and returns Gini to its no-attack baseline.

One property of the amended design deserves emphasis because it exceeded the designer's target. Running the final architecture at headline calibration with and without the composed attack, the attacked economy produces *higher* GDP and a *higher* universal dividend than the unattacked one ($954T and $229K against $875T and $159K), at the cost of 17 percent of twin-income independence and about two Gini points. The mechanism is the exposure tax itself: capital deployed in the attack becomes taxable exposure routed to the fund, and attack financing pushed into safe harbors becomes venture financing. The design does not merely resist its strongest known attack; it monetizes it. We claim this antifragility only against the discovered attack set: a genuinely novel exploit class would sit outside it (Section 7).

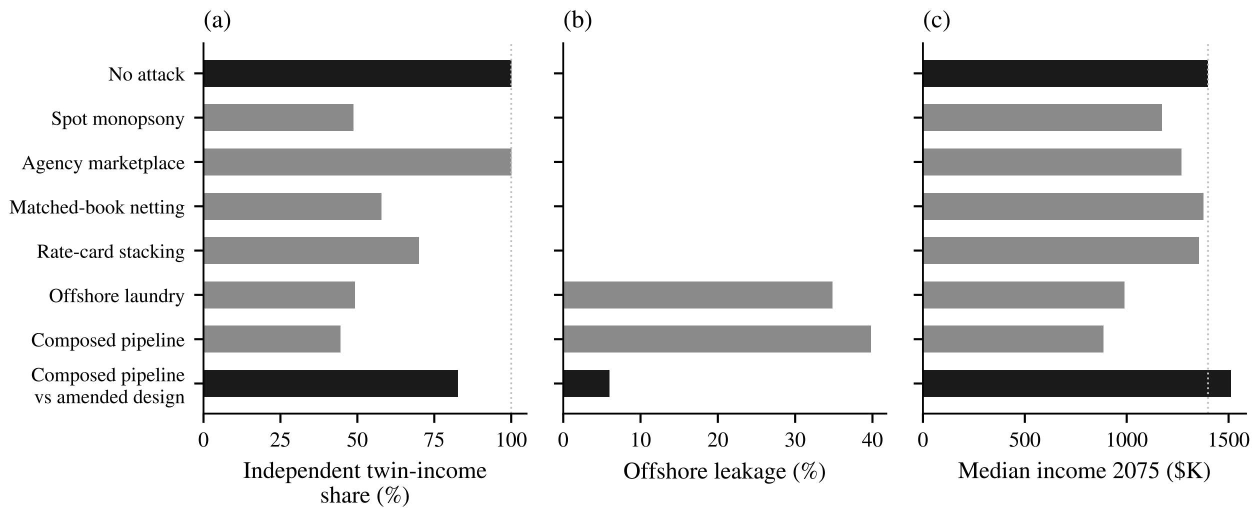

Figure 6: Round-three exploits versus the final design. Gray bars: attacks against the pre-amendment rules; black bars: no-attack baseline (top) and the amended design under the composed attack (bottom). The amendments restore twin-income independence to 83 percent, collapse offshore leakage from 40 to 6 percent, and recover median income.

6.3 Transition stress

The architecture's binding constraint is its transition window, when wages erode before the new channels mature. Under adverse pacing (automation arriving seven years early, twin adoption lagging), households below the floor peak at 7.3 percent in the early 2030s against 4.6 percent at baseline pacing, recover to zero by 2045, and the fund balance never goes negative (Figure 7). Uncomfortable but unbroken, conditional on the licensing and attestation infrastructure deploying on schedule. The two-sector Baumol extension adds a permanent honest caveat reported in Section 6.4.

6.4 Real income under Baumol dynamics

A two-sector extension with deflating goods and wage-indexed human services (Baumol, 1967) disciplines the paper's own real-income language. Against a basket frozen at 2025 weights, median real income grows 3.3-fold by 2075; against the evolving-weight basket that tracks actual consumption shares as services claim a rising share of spending, it grows 1.5-fold. Both numbers are true against their own basket; the difference is Baumol's cost

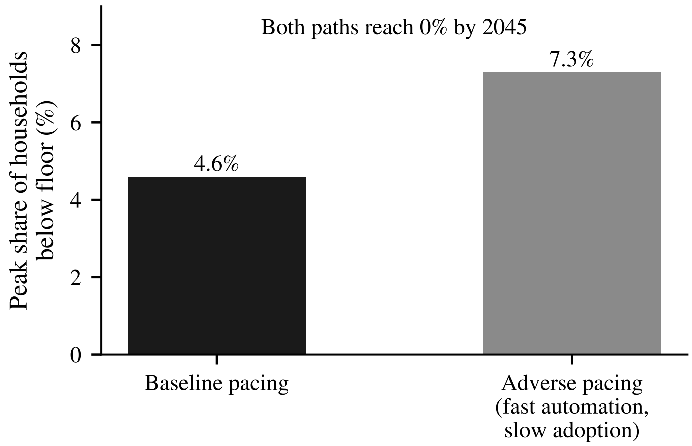

Figure 7: Transition stress test. Peak share of households below the floor under baseline and adverse pacing (automation seven years early, twin adoption lagging). Both paths recover to zero by 2045; the fund balance never goes negative.

disease honestly measured, and we report both throughout rather than choosing the flattering one. Two design consequences followed from this test: the floor is basket-indexed (a GDP-indexed floor loses roughly 8 percent real by 2075), and the basket weights are mechanized on transaction data because frozen weights eventually misstate what a floor is for.

7 Limitations

Six limitations bound the results. *Identification of κ.* The growth-advantage claim is conditional on cross-sectional identification of the demand-breadth feedback; within-country evidence cannot reject zero, and the required extrapolation (a 41-Gini-point counterfactual against 5 to 15 points of identifying variation) is severe. The comparative result degrades gracefully rather than flipping: at $\kappa = 0$ the architectures tie on growth and the claim reduces to distributional dominance. *Scenario dependence.* All levels are conditional on a near-complete automation trajectory; the model says nothing about whether that trajectory occurs, only about institutional comparisons along it. *Closed economy.* Offshore exit is modeled as an attack channel, not as open-economy general equilibrium; migration, trade, and international coordination are future work. *Attack-space bounds.* Three rounds of five personas plus a best-response search cover the exploits fifteen persona-passes could find; a novel exploit class would sit outside them. Known residuals are carried openly: the innovation-guard baseline is itself measurable and therefore a capture target; sovereign and state-owned AI fleets sit outside the tax perimeter; and the design's voting-adjacent machinery resists vote markets but not household-level coercion. *Attack parameterization.* Encoded attacks use structured estimates of take rates and reach derived from the red teams' qualitative findings, not calibrated data; the model verifies relative mechanism performance under attack, not absolute attack magnitudes. *Long-run institutional drift.* The model treats the definitions and enforcement institutions as constraints rather than dynamic players; multi-decade regulatory capture is flagged by the red team as the residual risk any paper design faces.

8 Conclusion

In an economy where machines generate most income, the assignment of that income is an institutional choice, and this paper's results say the choice has growth consequences. Across every tested strength of the demand-breadth feedback, an architecture that distributes AI-era earnings capacity universally out-grows one that lets machine earnings compound on existing wealth, while holding inequality near present-day levels instead of well above them. The mechanism is Schumpeterian rather than redistributive: broad purchasing power enlarges the prize for innovation, and universal capital access enlarges the set of people who can pursue it. The feedback parameter this rides on is empirically anchored at 0.37 by two independent routes, with the disclosed weakness that within-country evidence cannot reject zero; at zero, the growth claim reduces to distributional dominance at growth parity.

The paper's methodological contribution is independent of its policy conclusion: mechanism designs can be stress-tested the way security systems are, by structured adversaries whose exploits are encoded in the evaluating model before patches are accepted, under a pre-registered innovation guard that fails any design purchasing equality with dynamism. The license architecture survived three such rounds, and its strongest composed attack now raises the public dividend. Both the code and the full adversarial record are open; the standing invitation is to break either the mechanism or the anchoring.

Replication Package

The replication package accompanying this paper contains: the agent-based model code (versions 2 through 7, Python, with stage flags reproducing every table and figure); the cross-country regression data panel, construction code, and results (World Bank WDI, 2021 vintage); the three red-team council reports and the mechanism specification with all amendments; and model verification reports for each version. All figures in this paper are generated from the archived simulation outputs by the included scripts. Materials are available with the paper and from the author.

References

Daron Acemoglu and Joshua Linn. Market size in innovation: Theory and evidence from the pharmaceutical industry. *Quarterly Journal of Economics*, 119(3):1049–1090, 2004.

Daron Acemoglu and Pascual Restrepo. The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review*, 108(6):1488–1542, 2018. doi: 10.1257/aer.20160696.

Daron Acemoglu and Pascual Restrepo. Robots and jobs: Evidence from US labor markets. *Journal of Political Economy*, 128(6):2188–2244, 2020. doi: 10.1086/705716.

Philippe Aghion and Peter Howitt. A model of growth through creative destruction. *Econometrica*, 60(2):323–351, 1992.

David H. Autor. Why are there still so many jobs? The history and future of workplace automation. *Journal of Economic Perspectives*, 29(3):3–30, 2015. doi: 10.1257/jep.29.3.3.

Abhijit V. Banerjee and Esther Duflo. Inequality and growth: What can the data say? *Journal of Economic Growth*, 8(3):267–299, 2003. doi: 10.1023/A:1026205114860.

Robert J. Barro. Inequality and growth in a panel of countries. *Journal of Economic Growth*, 5(1):5–32, 2000. doi: 10.1023/A:1009850119329.

Emanuel F. Barros. Clutch economics: A post-job economy built on digital twin income, universal ownership, and consumer-driven growth. Working paper, v6. Companion materials: agent-based model code (v2–v7), three red-team council reports, mechanism design specifications, and the κ-anchoring report with cross-country regression data and code, 2026.

William J. Baumol. Macroeconomics of unbalanced growth: The anatomy of urban crisis. *American Economic Review*, 57(3):415–426, 1967.

Andrew Berg, Jonathan D. Ostry, Charalambos G. Tsangarides, and Yorbol Yakhshilikov. Redistribution, inequality, and growth: New evidence. *Journal of Economic Growth*, 23 (3):259–305, 2018. doi: 10.1007/s10887-017-9150-2.

Markus Brueckner and Daniel Lederman. Inequality and economic growth: The role of initial income. *Journal of Economic Growth*, 23(3):341–366, 2018. doi: 10.1007/s10887-018-9156-4.

Christopher D. Carroll, Jiri Slacalek, Kiichi Tokuoka, and Matthew N. White. The distribution of wealth and the marginal propensity to consume. *Quantitative Economics*, 8(3): 977–1020, 2017. doi: 10.3982/QE694.

Federico Cingano. Trends in income inequality and its impact on economic growth. OECD Social, Employment and Migration Working Papers 163, OECD Publishing, Paris, 2014.

Arnaud Costinot, Dave Donaldson, Margaret Kyle, and Heidi Williams. The more we die, the more we sell? A simple test of the home-market effect. *Quarterly Journal of Economics*, 134(2):843–894, 2019. doi: 10.1093/qje/qjz003.

Pierre Dubois, Olivier de Mouzon, Fiona Scott Morton, and Paul Seabright. Market size and pharmaceutical innovation. *RAND Journal of Economics*, 46(4):844–871, 2015. doi: 10.1111/1756-2171.12113.

J. Doyne Farmer and Duncan Foley. The economy needs agent-based modelling. *Nature*, 460(7256):685–686, 2009. doi: 10.1038/460685a.

Irving Fisher. *The Purchasing Power of Money: Its Determination and Relation to Credit, Interest and Crises*. The Macmillan Company, New York, 1911.

Reto Foellmi and Josef Zweimüller. Income distribution and demand-induced innovations. *Review of Economic Studies*, 73(4):941–960, 2006.

Kristin J. Forbes. A reassessment of the relationship between inequality and growth. *American Economic Review*, 90(4):869–887, 2000. doi: 10.1257/aer.90.4.869.

Carl Benedikt Frey and Michael A. Osborne. The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting and Social Change*, 114:254–280, 2017.

Daniel Halter, Manuel Oechslin, and Josef Zweimüller. Inequality and growth: The neglected time dimension. *Journal of Economic Growth*, 19(1):81–104, 2014. doi: 10.1007/s10887-013-9099-8.

Tullio Jappelli and Luigi Pistaferri. Fiscal policy and MPC heterogeneity. *American Economic Journal: Macroeconomics*, 6(4):107–136, 2014. doi: 10.1257/mac.6.4.107.

Xavier Jaravel. The unequal gains from product innovations: Evidence from the U.S. retail sector. *Quarterly Journal of Economics*, 134(2):715–783, 2019.

Damon Jones and Ioana Marinescu. 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, 2022. doi: 10.1257/pol.20190299.

Greg Kaplan and Giovanni L. Violante. A model of the consumption response to fiscal stimulus payments. *Econometrica*, 82(4):1199–1239, 2014. doi: 10.3982/ECTA10528.

Anton Korinek and Joseph E. Stiglitz. Artificial intelligence and its implications for income distribution and unemployment. In Ajay Agrawal, Joshua Gans, and Avi Goldfarb, editors, *The Economics of Artificial Intelligence: An Agenda*, pages 349–390. University of Chicago Press, Chicago, 2019. Also NBER Working Paper No. 24174 (December 2017).

Michael Kremer, Jonathan Levin, and Christopher M. Snyder. Advance market commitments: Insights from theory and experience. *AEA Papers and Proceedings*, 110:269–273, 2020. doi: 10.1257/pandp.20201017.

Kiminori Matsuyama. The rise of mass consumption societies. *Journal of Political Economy*, 110(5):1035–1070, 2002. doi: 10.1086/341873.

Benjamin Moll, Lukasz Rachel, and Pascual Restrepo. Uneven growth: Automation's impact on income and wealth inequality. *Econometrica*, 90(6):2645–2683, 2022. doi: 10.3982/ECTA19417.

Kevin M. Murphy, Andrei Shleifer, and Robert W. Vishny. Income distribution, market size, and industrialization. *Quarterly Journal of Economics*, 104(3):537–564, 1989. doi: 10.2307/2937810.

Pedro Cunha Neves, Óscar Afonso, and Sandra Tavares Silva. A meta-analytic reassessment of the effects of inequality on growth. *World Development*, 78:386–400, 2016.

Jonathan D. Ostry, Andrew Berg, and Charalambos G. Tsangarides. Redistribution, inequality, and growth. IMF Staff Discussion Note 14/02, International Monetary Fund, Washington, DC, 2014.

Jacob Schmookler. *Invention and Economic Growth*. Harvard University Press, Cambridge, MA, 1966.

Joseph A. Schumpeter. *Capitalism, Socialism and Democracy*. Harper & Brothers, New York, 1942.

Robert M. Solow. A contribution to the theory of economic growth. *Quarterly Journal of Economics*, 70(1):65–94, 1956.

Leigh Tesfatsion. Agent-based computational economics: A constructive approach to economic theory. In Leigh Tesfatsion and Kenneth L. Judd, editors, *Handbook of Computational Economics*, volume 2, chapter 16, pages 831–880. North-Holland/Elsevier, Amsterdam, 2006.

World Bank. World development indicators. Washington, DC: The World Bank, 2021. Database, 2021 vintage. `https://databank.worldbank.org/source/world-development-indicators`.

Josef Zweimüller. Schumpeterian entrepreneurs meet Engel's law: The impact of inequality on innovation-driven growth. *Journal of Economic Growth*, 5(2):185–206, 2000. doi: 10.1023/A:1009889321237.

A Speculative Extensions

This appendix is deliberately separated from the main text. Everything in it is scenario extrapolation under assumptions substantially stronger than those of Sections 2 to 6, and none of the paper's claims depend on it. It is included because the companion working paper (Barros, 2026) develops these scenarios in full and readers may wish to see the bridge.

A.1 Macroeconomic envelope scenarios to 2075

The companion paper brackets long-horizon output with growth accounting in normalized (index) form, $Y_{2075}/Y_{2025} = (K_{2075}/K_{2025})^{\alpha}(A_{2075}L_{2075}/A_{2025}L_{2025})^{1-\alpha}$, under a 2025 U.S. baseline (Y = \$25T, K = \$50T, L = 150M) (Solow, 1956). With capital scaled 20-fold, L = 25M voluntary entrepreneurs, and a terminal labor-side share of 0.9 ($\alpha = 0.1$), a productivity multiplier of 50 yields a 9.1-fold expansion (\$227T) and a multiplier of 117 yields 19.5-fold (\$487T). Because much of the automation literature expects the *capital* share to rise instead (Acemoglu and Restrepo, 2018), the same arithmetic at $\alpha = 0.9$ with only the conservative multiplier gives approximately \$458T: the two views land in the same corridor by different routes, and Section 4.3 tests the dispute at mechanism level. Monetary consistency is enforced exactly: every scenario satisfies $g_M + g_V = \pi + g_Y$ (Fisher, 1911), with velocity roughly doubling over the horizon and inflation at target.

A.2 The agent-based headline at anchored κ

At the anchored $\kappa = 0.37$, the final amended architecture produces the companion paper's headline scenario: real GDP of \$875T in 2075 (35-fold), median household income of \$1.69M nominal (23-fold; 1.5-fold real against evolving basket weights, 3.3-fold against frozen weights, per Section 6.4), income Gini 0.437, top-1 percent wealth share 18.9 percent (a smaller share of a much larger total, roughly 25 times richer absolutely), universal dividend \$159K, no household below the floor, and the innovation guard passing at a launch rate of 0.052 against a threshold of 0.042. Under the composed round-three attack the same design yields \$954T and a \$229K dividend (Section 6.2). We repeat the epistemic status: these are model outputs conditional on near-complete automation, anchored feedback, and the architecture as specified, not forecasts.

A.3 A 2150 extension with non-mortal agents

If life extension or machine sentience produces economically immortal agents, three assumptions behind most long-run distributional analysis fail: inheritance stops recycling fortunes, compounding horizons become unbounded, and population only grows. Extending the model to 2150 with no deaths after 2060, 0.8 percent annual population growth, and the socialization schedule escalating for century-scale fortunes, concentration remains bounded: the top-1 percent wealth share reaches 27.6 percent against 26.0 percent in the mortal baseline, no household falls below the floor, and median income remains stable while population doubles. Income Gini drifts to 0.51 over the additional 75 years, suggesting the escalation schedule should engage earlier. The bounding mechanism is the same maturity-linked rent decay that maximizes growth in Section 4.2: rules that reward continuous creation rather than accumulated position scale unchanged from mortal to non-mortal economies. The companion paper develops the institutional details, including the continuity rule that prevents replication from multiplying claims.

A.4 Relation to the companion working paper

The companion paper (Barros, 2026) presents the full institutional design (license law, attestation, fund mandate, definitions court, priority prizes), the complete red-team reports, a transition roadmap beginning under current law, and the long-horizon abundance scenarios. This paper extracts and defends its analytical core; where the two documents differ in emphasis, this paper's more conservative statements govern.