The Peer-Review Paper, Explained for Everyone

A plain-language guide to “Demand Breadth and Innovation-Driven Growth in a Highly Automated Economy: An Agent-Based Analysis” by Emanuel F. Barros (2026)

This guide assumes no background in economics or technology. It explains what the paper claims, how every part of the evidence works, what every symbol means, and who every cited researcher is. Reading time: about 45 minutes. A separate primer covers the broader Clutch Economics framework paper; this one covers the version written for professional economists.

Why two papers exist

Emanuel Barros has written two related papers.

The first, **“Clutch Economics” (the framework paper)**, is the full vision: Digital Twins, the citizen fund, the income floor, the transition roadmap, and long-range scenarios reaching to 2075 and beyond. It is written to be read by anyone.

The second, the one this guide explains, is the **peer-review paper**. It takes the single most defensible scientific claim inside the framework, strips away everything speculative, and presents it in the strict format professional economists use to judge each other’s work. Its bold cousin’s boldest material still appears, but quarantined in a clearly labeled appendix, the academic equivalent of a “speculation ahead” road sign.

What “peer review” means

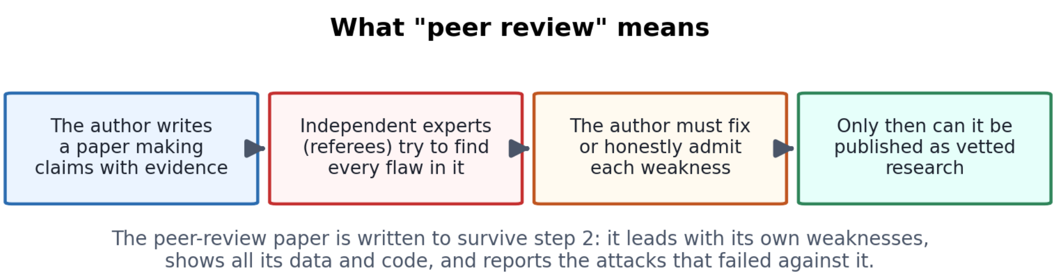

What peer review means

Academic journals do not publish papers because they sound convincing. They send each submission to independent experts (referees) whose job is to find every flaw: hidden assumptions, statistical tricks, missing evidence, wishful conclusions. Most submissions are rejected. Surviving peer review is how a claim graduates from “interesting idea” to “vetted research.”

So the peer-review paper is built the way you build for hostile inspection: every claim sized to its evidence, every weakness stated before a referee can discover it, all data and computer code published so anyone can re-run everything. Understanding that mindset explains almost every difference from the framework paper.

A map of the paper

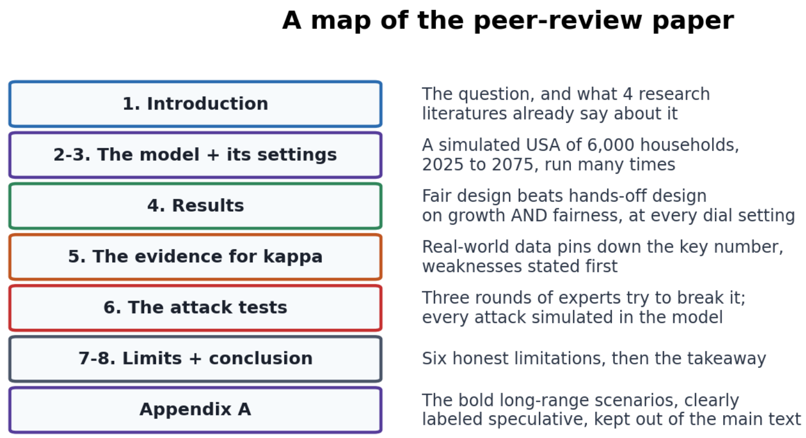

A map of the peer-review paper

---

1. The question the paper asks

Here is the setup in one breath. Suppose AI and robots eventually do most work. Then most income will be earned by machines, and society must DECIDE who receives it, because the automatic answer (whoever owns the machines) is just one possible rule among many. That makes the distribution of income a design choice for the first time in history.

The paper's question: **does that choice affect how fast the economy grows?** Specifically: if machine income is spread broadly across everyone, does the economy grow faster or slower than if it concentrates with the owners of the machines?

The standard instinct says slower: spreading money around dulls incentives, concentration builds capital. The paper's answer, from simulation plus real-world data, is the opposite: **the broad-distribution design grows FASTER, at every tested strength of the key feedback, while staying far less unequal.** Growth and fairness stop being a trade-off in an automated economy; they become the same mechanism seen from two sides.

Why anyone should take that seriously: the question is old and respectable

The paper opens by rooting itself in four established research traditions. This positioning matters for referees, who reflexively distrust claims that pretend to come from nowhere.

Tradition 1: Automation economics. Real studies of what robots and AI actually do to jobs and incomes.

Who's who here? Daron Acemoglu (MIT, among the most cited economists alive) and Pascual Restrepo (Boston University) wrote "Robots and Jobs" (2020), measuring how U.S. regions that got industrial robots lost jobs and wages, and "The Race between Man and Machine" (2018), the leading theory of automation versus new-task creation. David Autor (MIT) wrote "Why Are There Still So Many Jobs?" (2015), the classic explanation of why past automation waves did not end work (machines complement humans on some tasks), a caution the paper takes seriously. Carl Benedikt Frey and Michael Osborne (Oxford, 2017) produced the famous estimate that about 47 percent of U.S. jobs were at high risk of computerization. Anton Korinek (Virginia) and Joseph Stiglitz (Columbia, Nobel laureate) wrote the standard chapter on AI's implications for income distribution (2019). Benjamin Moll (LSE), Lukasz Rachel (UCL), and Pascual Restrepo published "Uneven Growth" (2022) in the field's top theory journal, showing automation channels gains to wealth owners. The paper takes this literature's displacement scenario as its starting premise.

Tradition 2: Market size and innovation. Evidence that inventors go where the customers are.

Who's who here? Jacob Schmookler's *Invention and Economic Growth* (1966) first documented that invention follows demand. Acemoglu and Joshua Linn (2004) showed pharmaceutical companies develop more drugs for bigger patient markets. Pierre Dubois and colleagues (2015) measured the same effect with a smaller elasticity. Arnaud Costinot, Dave Donaldson, Margaret Kyle, and Heidi Williams (2019) confirmed it using disease patterns across countries ("The More We Die, The More We Sell?"). Xavier Jaravel (LSE, 2019) showed new retail products flow toward customer groups whose spending power is growing. This is the paper's causal bedrock: demand size demonstrably pulls innovation.

Tradition 3: Income distribution and demand. Theory of how the SPREAD of income (not just its total) shapes what gets invented and produced.

Who's who here? Kevin Murphy, Andrei Shleifer, and Robert Vishny (1989) wrote two famous papers arguing industrialization needs a middle class big enough to buy factory goods ("Income Distribution, Market Size, and Industrialization" and "Industrialization and the Big Push"). Reto Foellmi and Josef Zweimüller (2006) built the modern theory of how inequality shapes demand-driven innovation; Zweimüller's earlier "Schumpeterian Entrepreneurs Meet Engel's Law" (2000) connects rich-versus-poor spending patterns to growth. Kiminori Matsuyama (2002) modeled "The Rise of Mass Consumption Societies": productivity gains spread purchasing power, which enlarges markets, which drives more productivity. The paper's demand-breadth mechanism is these ideas, updated for an automated economy.

Tradition 4: Who actually spends an extra dollar. Households differ: give a dollar to a family living tight and they spend most of it; give it to a wealthy family and they mostly save it. Economists call this the **marginal propensity to consume (MPC)**.

> **Who's who here?** Tullio Jappelli and Luigi Pistaferri (2014) documented MPC differences across households. Christopher Carroll and co-authors (2017) showed how wealth concentration lowers aggregate spending response. Greg Kaplan and Giovanni Violante (2014) built the standard model of why even middle-class families often spend windfalls. Together: moving income toward the broad middle demonstrably raises total spending, one ingredient of the paper's feedback.

There is also a fifth, more contested tradition, the cross-country inequality-and-growth literature, which the paper uses carefully in Section 5 of this guide.

---

2. The laboratory: a simulated economy

Since no one can run a 50-year experiment on a real country, the paper builds what economists call an **agent-based model**: a simulated economy of 6,000 households (statistically scaled to represent 250 million U.S. adults), each with its own starting wealth (matched to the real, highly concentrated U.S. wealth distribution), its own inventive talent (heavy-tailed: a few people very talented, most people modestly so, like reality), making decisions year by year from 2025 to 2075. Every configuration is run repeatedly with different random draws, and the paper reports median outcomes, so nothing rests on a lucky roll.

> **Who's who here?** Agent-based modeling is an established method: Leigh Tesfatsion (Iowa State) wrote the standard handbook treatment (2006), and J. Doyne Farmer (Oxford) and Duncan Foley argued in *Nature* (2009) that economics needs exactly this kind of modeling for questions where unequal individuals interact, questions that single-equation models average away. The distribution of income across thousands of different households is precisely such a question, which is why this method fits this paper.

The invention decision: the model's beating heart

The single most important design choice: **invention is a decision made inside the model, not an assumption fed into it.** Each simulated household weighs whether launching a venture is worth it:

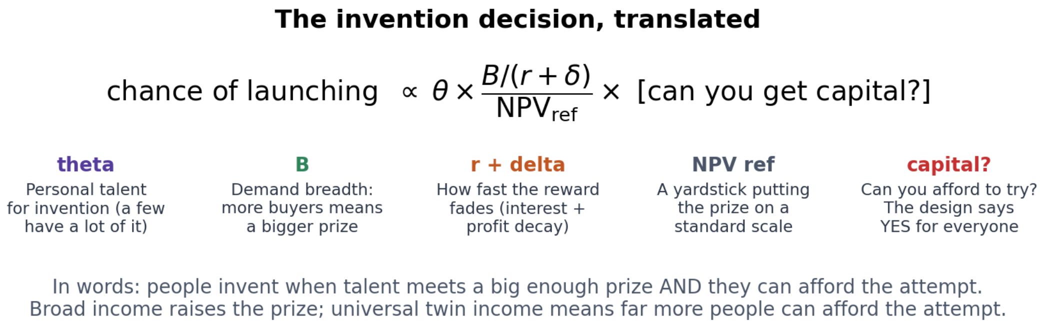

The invention decision, translated

In words: your chance of launching a venture rises with your talent (**theta**), rises with the size of the prize (which grows when demand is broad, **B**, and shrinks the faster profits decay, **delta**, against the interest rate **r**, all measured against a standard yardstick called **NPV_ref**), and requires that you can get capital at all. That last piece is where institutions enter: in the broad-distribution design, everyone has income and can finance an attempt; in the hands-off benchmark, only the wealthy can.

Because invention is a choice, the model can PUNISH the designer: any rule that spoils the inventor's reward shows up as collapsing venture creation. The paper formalizes this with the **Innovation Guard**: any configuration where venture launches fall below 90 percent of a completely unregulated economy's rate is declared a failure, no matter how pretty its other numbers.

The Innovation Guard

The growth equation

Each year, the simulated economy grows by:

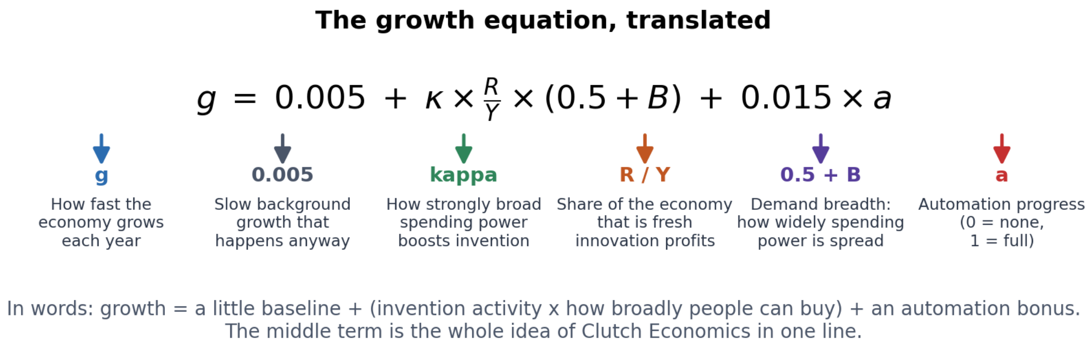

The growth equation, translated

The middle term is the paper's thesis in symbols: growth from innovation scales with how much fresh innovation profit is in the economy (**R/Y**) times how broadly people can buy (**B**, which is 1 minus the Gini inequality measure), with the strength of that link set by one number: **kappa**. Everything the paper claims about growth depends on kappa being above zero, which is why the paper devotes a full section to measuring it (Section 5 of this guide).

The two contenders

The model runs identical people, identical technology, identical random events, under two rule-sets:

The license architecture (the proposal): every adult holds one non-transferable license to an AI agent earning income on their behalf; a citizen fund pays an equal dividend out of commoditized profits and a share of large passive returns; innovators keep decades-long rewards; a guaranteed floor sits underneath. (The framework paper calls this the Digital Twin design; the peer-review paper deliberately uses the drier name "license architecture.")

Laissez-faire (the benchmark): machine earnings flow to whoever owns machine capital, in proportion to wealth. No license, no dividend, no floor. Importantly, this benchmark is NOT rigged to fail: it enjoys the same technology, the same innovation system, and the same profit protections. It differs in exactly one thing: who receives the machines' earnings.

---

3. The results, one experiment at a time

Experiment 1: the headline

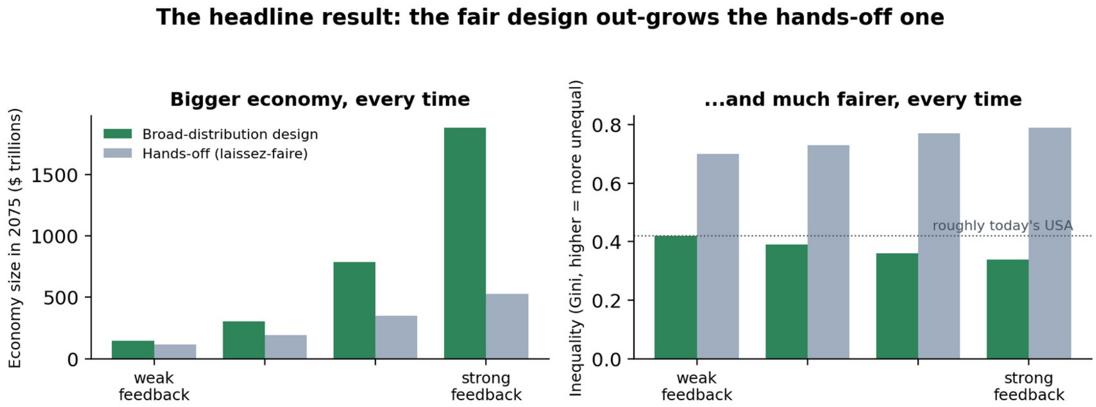

The headline result

At every tested strength of the demand feedback (kappa from 0.10 to 0.50), the license architecture produces a larger 2075 economy than laissez-faire, and the gap widens as feedback strengthens: 144 versus 117 trillion dollars at the weakest setting, 786 versus 351 trillion at the middle, 1,881 versus 527 trillion at the strongest. Inequality tells the mirror story: Gini around 0.34 to 0.42 (near today's) under the license design versus 0.70 to 0.79 (extreme concentration) under laissez-faire. Zero simulated households fall below the floor.

Why? The paper switches each mechanism off separately to find out. Two channels do the work: **bigger prizes** (broad demand makes each invention worth more) and **more players** (universal income means far more people can afford to attempt ventures). Concentration starves both. Notably, if kappa is set to zero, the two designs tie on growth, and the license design merely wins on fairness. The paper says this openly rather than burying it: the growth claim lives or dies with kappa.

Experiment 2: how long should inventors keep their profits?

Sweeping the profit-protection half-life from 2 years to 40 years: growth rises steeply with longer protection (69 trillion to 870 trillion) while inequality rises only gently (Gini 0.20 to 0.37). Lesson: the fairness work is done by the license and the fund, not by hurrying inventors' rewards away. The design therefore adopts patent-length (25 to 40 year) protection. The paper highlights that this OVERRULED the author's own initial preference for fast decay, which is worth noticing: a model that never contradicts its creator is usually a mirror, not an instrument.

Experiment 3: the strongest objection, met head-on

The most predictable expert objection: "Mainstream automation research (Acemoglu-Restrepo, Moll and co-authors) says machine OWNERS will capture a rising share of income. Your design leans on twin earnings, a kind of labor income. Automation will crush it." The paper runs exactly that world, with capital's share of output ramping to 60 percent. The result inverts the objection: the design performs BETTER (economy of 828 versus 635 trillion; median income 1.73 versus 1.25 million; Gini falling to 0.31), because swollen capital profits flow into the citizen fund and the dividend automatically replaces shrinking twin income. The fund is a built-in stabilizer for precisely the future the objectors predict. An adaptive variant that adds a direct citizen share of capital income does better still (983 trillion). The honest nuance, stated in the paper: under high capital shares, the dividend rather than the twin becomes the main income channel, so the argument shifts from "twins earn your living" to "your fund share earns your living."

Experiment 4: which knobs matter?

A "tornado analysis" wiggles every model parameter up and down by half its value and charts which movements change the outcome most. Result: **kappa dominates everything** (2075 output swings from 213 to 1,517 trillion across its range, roughly triple any other parameter's effect). The licensed pool share matters most for fairness. Nothing in the sweep breaks the Innovation Guard. This experiment is why the paper then goes hunting for kappa in real-world data: when your conclusion leans on one number, you owe the world a measurement, not an assumption.

---

4. Measuring kappa: the paper's most important and most honest section

Kappa is the strength of the link between demand breadth and innovation-driven growth. The paper pins it down two independent ways, and then, unusually, leads with the weakness of its own evidence.

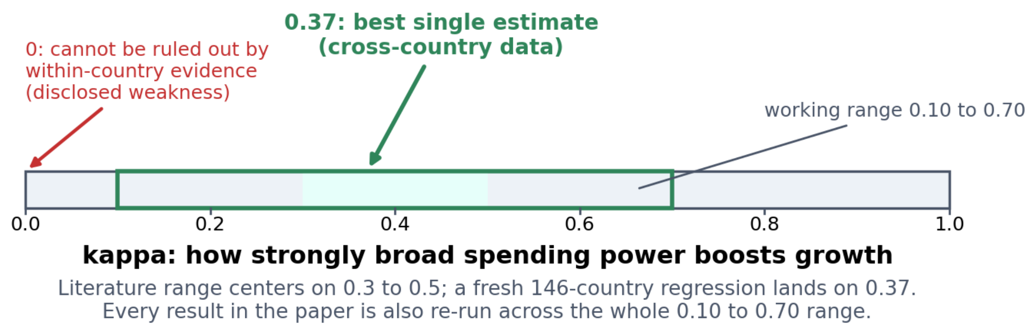

Where kappa's estimate comes from

Route 1: triangulating from existing research

Dozens of published studies measure things adjacent to kappa. Converted into kappa's units (the paper shows the conversion arithmetic):

- The **market-size studies** (Acemoglu-Linn, Costinot et al., Dubois et al., Jaravel; see Section 1) imply kappa between roughly 0.03 and 0.9, with the causal studies favoring the middle.
- The **household spending studies** (Jappelli-Pistaferri, Carroll et al., Kaplan-Violante) imply about 0.12 to 0.17 from the spending channel alone.
- The **cross-country inequality-growth studies** imply values up to 1 or 2, but the paper treats those as ceilings, not estimates, because they bundle every possible channel (education, instability, credit) together, while kappa is only the demand-breadth channel.

  Who's who in that third group? Federico Cingano's OECD study (2014) and the IMF team of Jonathan Ostry, Andrew Berg, and Charalambos Tsangarides (2014, extended in 2018 with Yorbol Yakhshilikov) found inequality associated with slower growth. Markus Brueckner and Daniel Lederman (2018) found the effect depends on a country's income level. Pedro Neves, Óscar Afonso, and Sandra Silva (2016) pooled all prior studies (a meta-analysis) and found a small average effect. And the field's famous disagreement: Kristin Forbes (2000) and Robert Barro (2000) showed that WITHIN one country over time, inequality can look growth-neutral or even positive, while comparisons ACROSS countries look negative; Abhijit Banerjee (Nobel laureate) and Esther Duflo (Nobel laureate) (2003) warned the relationship is not a straight line; Daniel Halter, Manuel Oechslin, and Josef Zweimüller (2014) showed the sign flips with time horizon. The paper inherits this messy literature honestly rather than cherry-picking its favorable half.

Bottom line from route 1: kappa plausibly sits between 0.1 and 1.0, most likely 0.3 to 0.5.

Route 2: a fresh measurement

The paper runs its own statistical analysis (a **regression**, the standard tool for measuring how one thing moves with another while holding other factors constant) on World Bank data covering up to 167 countries from 1990 to 2019. It asks: do countries where spending power is spread more broadly grow faster, after accounting for how rich they already are, how much they invest, and how educated they are?

Two ways of slicing the data give two different answers, and understanding why is the key to the whole section:

- **Comparing countries against each other over the long run** (the "between" comparison): clearly yes. Ten points more demand breadth associates with about 0.37 percentage points more yearly growth, which translates to kappa of about 0.37. The odds of seeing this pattern by pure chance are about 4 in 1,000 (that is what "p = 0.004" means). The finding survives restricting to rich countries and trimming outliers.
- **Tracking each country against its own past** (the "within" comparison): no detectable effect. Inequality moves too slowly inside one country, and the data is too noisy, for this test to see anything, including, possibly, a real effect. But a fair reading must allow that the effect could be zero.

An everyday analogy: comparing countries is like noticing that people who exercise more tend to be healthier (suggestive, but maybe healthy people just like exercise); tracking one person over time would be better evidence, but if your fitness tracker only logs data once a decade, it will detect nothing either way. The paper says, in effect: the suggestive comparison points to 0.37, the strict test is silent, and a skeptic is entitled to remain one.

Adopted value: kappa = 0.37, with every result re-run across 0.10 to 0.70. And the paper repeats its central honesty: at kappa = 0, the growth advantage disappears and only the fairness advantage remains. Nothing flips into reverse; the claim just shrinks.

(A technical footnote you can skip: the added-variable chart in the paper shows each of 146 countries as a dot, with the paper's measured slope drawn through them. The exact numbers: slope 0.037, standard error 0.013, p = 0.004. This guide's author re-computed that regression independently from the published data files and got identical numbers.)

---

5. The unusual chapter: robustness to strategic manipulation

Most policy papers end with "of course, people might game this." This paper instead hired the gamers. Three separate rounds of five adversarial expert personas (a loophole lawyer, a platform monopolist, a financial engineer, a political lobbyist, and a gray-market economist) attacked the design. Every attack was then **programmed into the simulation as a strategy** and run, so its damage is a number, not a fear. Only then were rules patched, and the next round attacked the patches. The paper argues this is itself a contribution to economics: mechanism designs can be security-tested like software, and publishing the full attack record invites the world to find what fifteen expert passes missed.

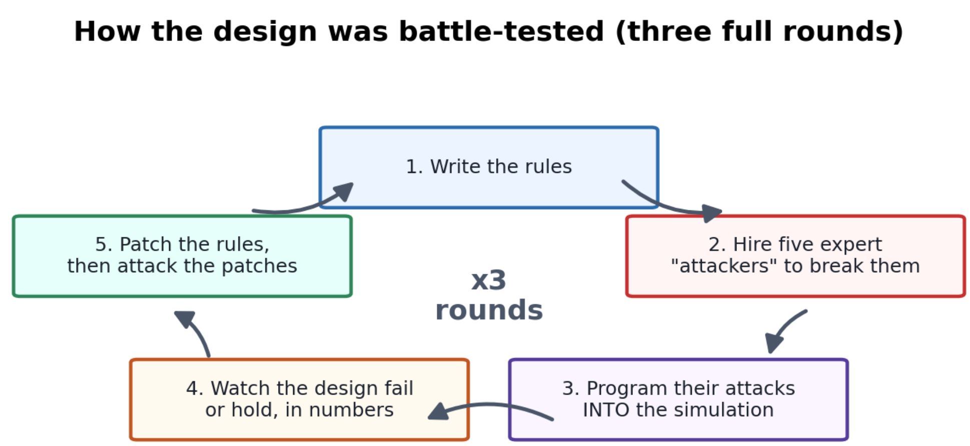

The red-team cycle

The three rounds in brief (the framework-paper primer tells this story in full; here is what the peer-review paper reports):

- **Round one** found that buying a citizen's future income forever, bundled with an "advisory" contract, effectively buys the citizen's twin. Simulated unpatched: advisors capture 36 percent of twin income, top-1 percent wealth hits 45 percent, and the economy is less than half its patched size. Patched (income pledges capped at 3 years and 35 percent, mass-advisory services structurally separated, verification of licensed uniqueness, ownership look-through): capture falls to 5 percent and the system holds. An automated worst-case search on top of the human attackers found only graceful degradation (Gini 0.49 versus 0.44 baseline; nobody below the floor).
- **Round two** rebuilt corporate fleets through revenue-based financing and exposed a prize add-on as a capture surface (the model showed the add-on bought zero extra innovation, so it was deleted). The surviving fix is the **economic-exposure rule**: any party whose total economic stake in a fleet's output tops 20 percent, through ANY instrument, is taxed as an owner. Not banned, taxed: the model separately showed outside financing HELPS innovation (locking capital-poor talent out of scale costs growth), so the rule targets disguised control, not capital itself.
- **Round three** attacked the exposure rule and found the project's worst single exploit: **offshore laundering** of twin output through unverified foreign middlemen. Simulated, it drained 35 percent of the twin income pool offshore and was the only attack in all three rounds to break the Innovation Guard. The final amendment (an export withholding tax at unverified borders, refunded when money verifiably returns onshore) collapses the leak from 40 percent to 6 percent.

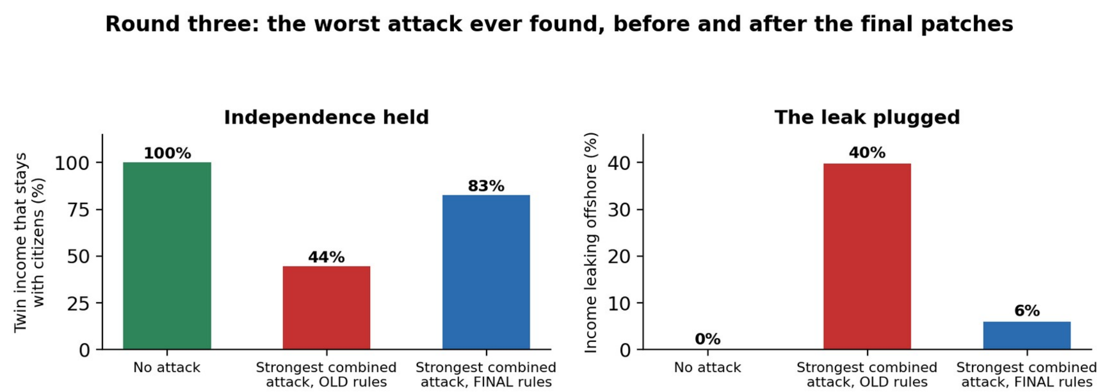

The worst attack, before and after

The closing twist: under the strongest COMBINED attack the rounds could compose, the finished design produces higher output and a higher citizen dividend than with no attack at all (954 versus 875 trillion; 229 versus 159 thousand dollars), because the exposure tax converts attackers' capital into fund revenue and squeezes their financing into legitimate venture funding. The paper is careful about the claim's edges: independence of twin income dips to 83 percent, inequality ticks up about two Gini points, and the property is claimed only against KNOWN attacks.

Two further robustness results round out the chapter. A **transition stress test** (automation arriving seven years early, adoption lagging) peaks at 7.3 percent of households below the floor in the early 2030s, recovering to zero by 2045, with the fund never going broke. And a **Baumol two-sector test** (machines make goods cheap, human services stay dear; from William Baumol's classic 1967 paper) disciplines the paper's own real-income language: median real income grows 3.3-fold against a frozen 2025 shopping basket but 1.5-fold against a realistically evolving one, and the paper reports both numbers everywhere rather than choosing the flattering one.

---

6. The quarantine zone: Appendix A

Everything spectacular lives in a clearly labeled appendix that opens by telling you not to lean on it. This is deliberate scientific hygiene: the main text's claims are meant to survive hostile review; the appendix shows where the framework goes if you extend the assumptions, and points to the framework paper for the full treatment.

Inside the quarantine: the long-range 2075 scenarios (a 227 to 487 trillion dollar envelope from growth accounting, and the 875 trillion headline from the simulation at measured kappa, with median household income of 1.69 million); and the strangest extension, a **2150 economy with immortal participants** (if life extension or sentient machines mean some economic actors never die, does wealth compound into permanent aristocracy? In the model, no: the same maturity-based decay of innovation profits that maximizes growth also forces even immortal fortunes to keep creating, and top-1 percent wealth stays near 27.6 percent versus 26.0 in the mortal baseline).

---

7. The paper's own list of weaknesses

Section 7 of the paper lists six limitations without hedging. In plain words:

1. **Kappa's evidence is cross-sectional.** Country-versus-country comparisons support 0.37; country-versus-its-own-past cannot confirm it. A skeptic may hold kappa near zero, in which case the design merely ties on growth and wins on fairness.
2. **Everything is conditional on the automation scenario.** The paper compares institutions along an assumed trajectory; it does not prove the trajectory.
3. **The model is a closed single country.** Offshore behavior appears only as an attack channel; full international economics is future work.
4. **The attackers were thorough, not omniscient.** A genuinely novel exploit class could exist outside the fifteen persona-passes and the automated search. Known soft spots are listed openly (the guard's own baseline could be gamed; government-owned AI fleets sit outside the tax rules; ballots resist vote-buying but not household pressure).
5. **Attack sizes are informed estimates**, not calibrated measurements; the model verifies which rules beat which attacks, not exact real-world magnitudes.
6. **Slow institutional capture** (decades of lobbying) is designed against but not dynamically simulated.

---

8. How to read the paper's numbers responsibly

Three habits the paper itself models, worth adopting when you discuss it:

1. **Ratios over levels.** "The license design produced roughly twice laissez-faire's output" is robust across assumptions. "The economy will be 875 trillion dollars" is a scenario output, not a forecast, and the paper never claims otherwise.
2. **Both real-income numbers.** Median income grows 23-fold in dollars, 3.3-fold against a frozen basket, 1.5-fold against an evolving basket. All three are true statements about different questions; quoting only the biggest one would be spin.
3. **The kappa caveat travels with the headline.** Any honest summary of this paper includes "if the demand-breadth feedback is real, and cross-country evidence suggests it is."

---

9. Who's who: the complete annotated reference list

Every work cited in the paper, in alphabetical order: who, what, and why it is there.

- **Acemoglu & Linn (2004).** MIT economists; showed drug development follows market size. Causal bedrock for demand-pulls-innovation.
- **Acemoglu & Restrepo (2018).** Theory of automation versus new task creation in the American Economic Review. Frames the displacement scenario.
- **Acemoglu & Restrepo (2020).** The robots-and-jobs measurement in U.S. labor markets. Empirical ground truth for displacement.
- **Aghion & Howitt (1992).** Philippe Aghion and Peter Howitt built the standard mathematical model of growth through "creative destruction," innovators displacing incumbents. The paper's long innovation rewards are justified in this tradition.
- **Autor (2015).** Why automation has historically complemented humans; the standard caution against naive "jobs are doomed" claims.
- **Banerjee & Duflo (2003).** Nobel-winning development economists; showed inequality-growth data resists simple straight-line stories. Cited as a warning label on extrapolation.
- **Barro (2000).** Harvard macroeconomist; found inequality's growth effect differs between rich and poor countries, and between statistical methods. Half of the famous estimator disagreement the paper inherits.
- **Barros (2026).** The companion framework paper and its published code, red-team reports, specifications, and data.
- **Baumol (1967).** The "cost disease": human services get relatively dearer forever. Forces the paper's honest double reporting of real income.
- **Berg, Ostry, Tsangarides & Yakhshilikov (2018).** IMF economists; updated evidence that redistribution itself is not growth-harmful. Context for the fund's design.
- **Brueckner & Lederman (2018).** World Bank and Australian National University; inequality's growth effect depends on initial income. More context on the contested literature.

- **Brynjolfsson & McAfee (2014).** MIT; the standard popular-academic account of digital transformation of work.
- **Carroll, Slacalek, Tokuoka & White (2017).** Johns Hopkins and central-bank economists; measured how spending responses differ across the wealth distribution. Ingredient of the demand-breadth feedback.
- **Cingano (2014).** OECD; documented inequality associated with slower growth across member countries. Upper-bound input to kappa.
- **Costinot, Donaldson, Kyle & Williams (2019).** MIT and Tuck; used global disease differences to confirm home-market demand drives pharma innovation.
- **Dubois, de Mouzon, Scott Morton & Seabright (2015).** Toulouse and Yale; measured a smaller market-size elasticity for drugs. Keeps kappa's lower range honest.
- **Farmer & Foley (2009).** Complexity scientist and economist; the *Nature* manifesto for agent-based economics. Methodological justification.
- **Fisher (1911).** The money-flow identity every scenario respects.
- **Foellmi & Zweimüller (2006).** Swiss growth theorists; the modern model of inequality shaping demand-induced innovation. The paper's mechanism, formalized a generation earlier.
- **Forbes (2000).** MIT; found within-country inequality changes can associate with FASTER growth, the sharpest version of the estimator disagreement the paper discloses rather than hides.
- **Frey & Osborne (2017).** The famous 47-percent-of-jobs-at-risk estimate. Scenario premise.
- **Grossman & Helpman (1991).** Princeton and Harvard trade-and-growth theorists; the book-length treatment of innovation-driven growth. Theoretical scaffolding.
- **Halter, Oechslin & Zweimüller (2014).** Showed inequality helps growth short-run, hurts long-run. Explains why the paper's long-horizon between-country estimate is the relevant one.
- **Jappelli & Pistaferri (2014).** Documented MPC heterogeneity: modest households spend extra income; wealthy households park it.
- **Jaravel (2019).** LSE; product innovation flows to customer groups with growing purchasing power. The retail-sector twin of the pharma evidence.
- **Jones & Marinescu (2022).** The Alaska dividend study: permanent universal cash did not reduce employment. Real-world anchor for the dividend channel.
- **Kahneman & Tversky (1979).** Prospect theory, the foundation of behavioral economics. Cited for loss-aversion psychology.
- **Kaplan & Violante (2014).** The "wealthy hand-to-mouth" model: why even middle-class families spend windfalls. MPC ingredient.
- **Korinek & Stiglitz (2019).** The standard chapter on AI, inequality, and unemployment; frames the institutional-choice question the paper answers.
- **Kremer, Levin & Snyder (2020).** Nobel laureate Kremer and colleagues on advance market commitments: prizes work in the real world (vaccines). Precedent for the design's prize mechanism.

- **Matsuyama (2002).** Northwestern; the virtuous circle of mass consumption and productivity. Theoretical ancestor.
- **Moll, Rachel & Restrepo (2022).** Automation's gains flowing to wealth owners in the top theory journal. The stress-test target of Experiment 3.
- **Murphy, Shleifer & Vishny (1989, two papers).** Chicago and Harvard; industrialization needs a mass market. The oldest version of the paper's thesis.
- **Neves, Afonso & Silva (2016).** The meta-analysis (study of studies) of inequality and growth; small average effect. Keeps kappa's estimate conservative.
- **Ostry, Berg & Tsangarides (2014).** IMF; inequality harmful, redistribution benign. Contested but influential context.
- **Schmookler (1966).** Invention follows demand, documented across a century of patents. The tradition's founding text.
- **Schumpeter (1942).** Creative destruction; innovation needs prospective rewards. The design's incentive philosophy.
- **Solow (1956).** Growth accounting; used in the appendix's envelope arithmetic.
- **Tesfatsion (2006).** The handbook chapter defining agent-based computational economics. Methodological foundation.
- **World Bank WDI (2021 vintage).** The public country-statistics database behind the kappa regression.
- **Zweimüller (2000).** Inequality, Engel's law, and innovation-driven growth. Direct theoretical ancestor of demand breadth.

---

10. Small glossary

- **Added-variable plot:** a chart showing the relationship between two things after statistically removing other influences; the honest way to picture a regression.
- **Agent-based model:** a simulation of many individual decision-makers interacting under rules, run forward in time.
- **Between versus within estimates:** comparing different countries against each other, versus tracking one country against its own past. They disagree in this literature, and the paper discloses it.
- **Calibration:** setting a model's dials to match real-world measurements (U.S. wealth distribution, GDP, and so on).
- **Demand breadth (B):** 1 minus Gini; how widely spending power is spread.
- **Gini coefficient:** the 0-to-1 inequality scale (0 = everyone equal, 1 = one person has everything).
- **Innovation Guard:** the pass/fail rule: venture creation must stay at or above 90 percent of a free economy's rate.
- **Kappa:** the measured strength (best estimate 0.37) of the demand-breadth-to-growth feedback.
- **Laissez-faire:** the hands-off benchmark; machine earnings flow to machine owners.
- **Monte Carlo:** running a simulation many times with different random draws and reporting the middle result.

- **MPC (marginal propensity to consume):** the share of an extra dollar a household spends rather than saves.
- **p-value:** the probability of seeing a pattern at least this strong by pure chance; $p = 0.004$ means about 4 in 1,000.
- **Quasi-rent:** the temporary outsized profit an innovator earns before competition catches up.
- **Red team:** experts hired to attack your own design before real adversaries do.
- **Regression:** the standard statistical tool for measuring how one variable moves with another while holding others constant.
- **Replication package:** the published code and data letting anyone re-run every result.
- **Robustness:** whether a result survives changing the assumptions.
- **Tornado analysis:** wiggling every parameter to see which ones actually move the result, charted like a tornado's funnel.

This primer summarizes "Demand Breadth and Innovation-Driven Growth in a Highly Automated Economy: An Agent-Based Analysis" (July 2026) by Emanuel F. Barros. The paper's replication package (model code, regression data and code, adversarial-review reports) is published with it. For the broader framework, see the separate primer "Clutch Economics, Explained for Everyone."