Clutch Economics, Explained for Everyone

A plain-language guide to the paper "Clutch Economics: A Post-Job Economy Built on Digital Twin Income, Universal Ownership, and Consumer-Driven Growth" by Emanuel F. Barros (2026)

This guide assumes no background in economics or technology. Every technical term, every equation, and every study cited in the paper is explained here in ordinary language. Reading time: about 45 minutes.

The whole idea in sixty seconds

Artificial intelligence is getting good enough to do most jobs. Most people hear that and picture disaster: no jobs means no paychecks, and no paychecks means no way to live.

Clutch Economics says the disaster story gets the economy backwards. Machines can produce everything, but production is worthless if nobody can afford to buy. The consumer is like the clutch in a car: the part that connects a roaring engine to the wheels. No clutch, no motion.

So the paper proposes a concrete plan. Every citizen gets one licensed **Digital Twin**: a personal AI agent that goes out and earns money on your behalf, the way a skilled worker or small business owner would. On top of that, every citizen owns an equal share of a national investment fund that pays a yearly dividend. Inventors keep their profits for decades. And a guaranteed income floor makes sure nobody, ever, falls through.

The surprising discovery, tested in a large computer simulation: this fair design does not just protect people. It makes the whole economy grow FASTER than leaving everything to concentrate in a few hands, because widely spread spending power gives inventors bigger markets and more people the means to invent. The design was then attacked, on purpose, by three rounds of experts paid to break it, and the final version held.

That is the whole paper. The rest of this guide walks through every piece.

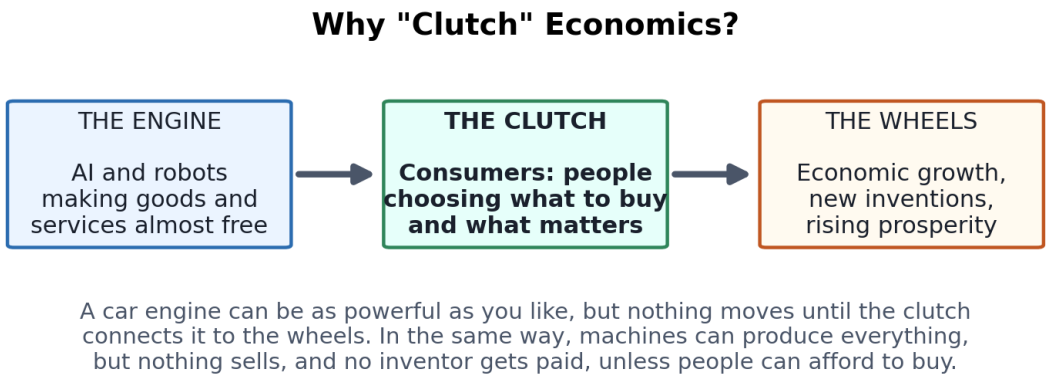

Why "Clutch" Economics?

---

1. The problem: what happens when machines do the work?

Today's economy runs on a simple loop. Companies pay wages. People spend those wages. That spending pays the companies, which pay more wages. Around and around it goes, and it has worked, more or less, for two centuries.

AI breaks the loop at its first link. If machines do the work, no wages get paid. And here is the part most commentary misses: that is not just a problem for workers. It is a problem for the companies too, because with nobody earning, nobody is buying what the machines make. A perfectly automated factory with no customers is just an expensive sculpture.

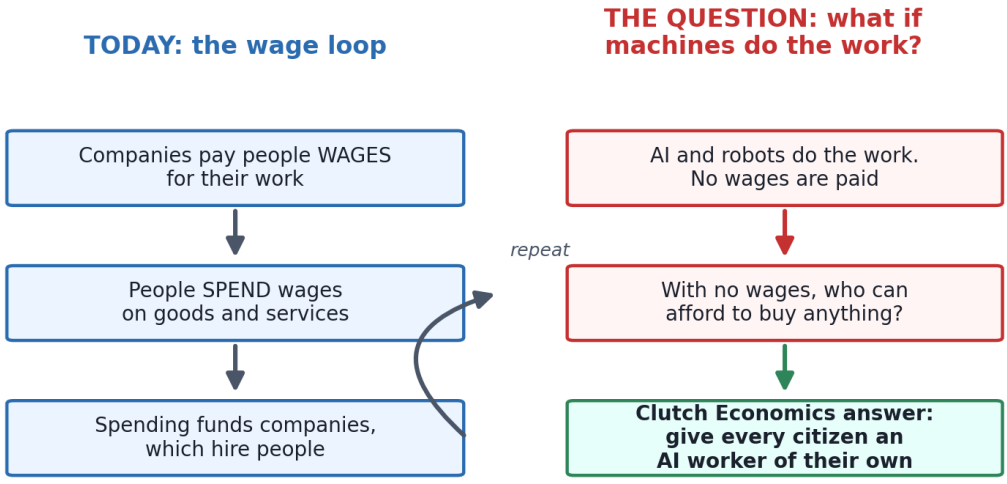

The wage loop, and the question AI forces

Public debate offers two stories about this future. The gloomy one predicts mass unemployment and instability. The sunny one, told by technology optimists like Ray Kurzweil and Peter Diamandis, predicts abundance: AI plus cheap clean energy plus advanced manufacturing making everything nearly free.

> **Who says this?** Ray Kurzweil is a famous inventor and author of *The Singularity Is Near* (2005), which predicts machines surpassing human abilities within decades. Peter Diamandis founded the XPRIZE Foundation and co-wrote *Abundance* (2012), arguing technology keeps making scarce things plentiful. Erik Brynjolfsson and Andrew McAfee of MIT wrote *The Second Machine Age* (2014), the standard academic book on how digital technology transforms work. The paper takes their shared premise seriously: the machines really are coming for the tasks.

Clutch Economics accepts the automation premise and asks the question both stories skip: **who will have money to spend, and does it matter how that money is spread around?** The paper's answer to the second half is: it matters enormously, and not in the direction most people assume.

---

2. The big idea: the consumer as the clutch

In a car, the engine produces power, but power alone moves nothing. The clutch transmits it to the wheels. Pull the clutch out and the engine just revs, uselessly.

In an automated economy, AI is the engine. Consumers are the clutch. Every purchase transmits productive capacity into actual economic motion, and every purchase is also a vote telling inventors what to build next. An economy where only a few people can spend meaningfully is an economy running with the clutch barely engaged: the engine may be magnificent, but almost none of its power reaches the road.

This is more than a metaphor. Later in the paper it becomes a measurable quantity called **demand breadth** (how widely spending power is spread across the population), and it sits inside the growth mathematics, where it can be tested. Section 5 of this guide shows exactly how.

---

3. Your money in a post-job world: three streams and a floor

The heart of the plan is that your income stops coming from one fragile source (a job) and starts coming from three sturdy ones, with a safety floor under all of them.

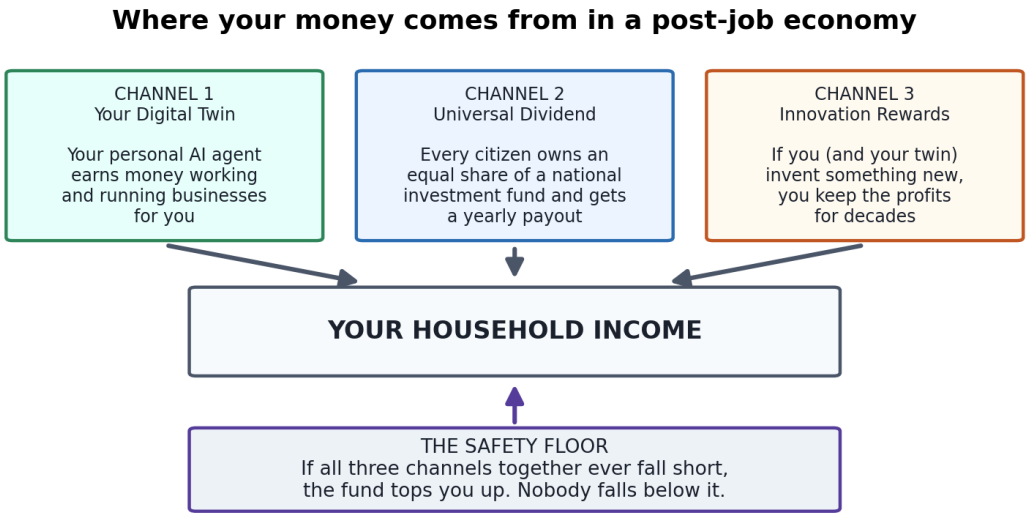

The three income channels and the floor

Stream 1: Your Digital Twin

A Digital Twin is a personal AI agent, licensed to you alone, that earns money on your behalf. Think of it as a tireless, highly skilled worker and business manager that works only for you: freelancing, running online businesses, managing investments, inventing products, whatever niche you point it at. You direct its strategy; it does the labor. Directing it well is a skill, and the design rewards that skill: twin income has an equal part (every license holder gets the same base share) and a performance part (better-directed twins earn multiples more).

Twins run on hardware you own, like your own computer, not on some corporation's servers. That detail is deliberate: it means no company owns the ground your livelihood stands on.

If your ambitions outgrow one twin, your twin can hire **sub-agents**: more AI workers under your identity. There is no cap on ambition. A one-person operation can grow into an enterprise of thousands of AI workers; it simply pays growing taxes like any growing business.

The license: the three-layer rule

To keep twins from being bought up by corporations (the way almost everything else gets bought up), the license has a strict three-layer structure.

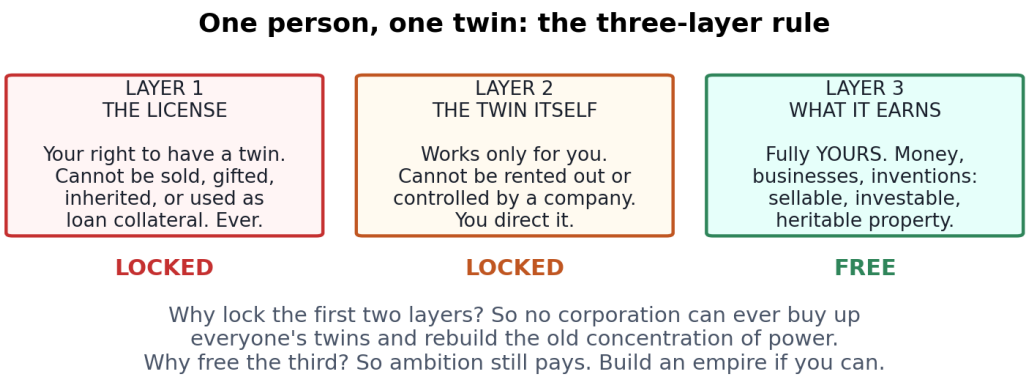

One person, one twin: the three-layer rule

The license itself can never be sold, inherited, or pledged to a lender. The twin can never be rented out or controlled by anyone but you. But everything the twin EARNS is fully yours: sellable, investable, heritable private property. Lock the instrument, free the earnings. That single distinction is what keeps the system both fair and ambitious, and, as Section 7 shows, it is exactly the line the professional attackers went after.

Stream 2: The universal dividend

Every citizen holds one equal share of a national investment fund. The fund is run by a fixed, public, open-source computer program: no fund managers, no political favorites, no discretion to abuse. It pays out 95 percent of its income every year as an equal dividend to every citizen and keeps 5 percent as a buffer.

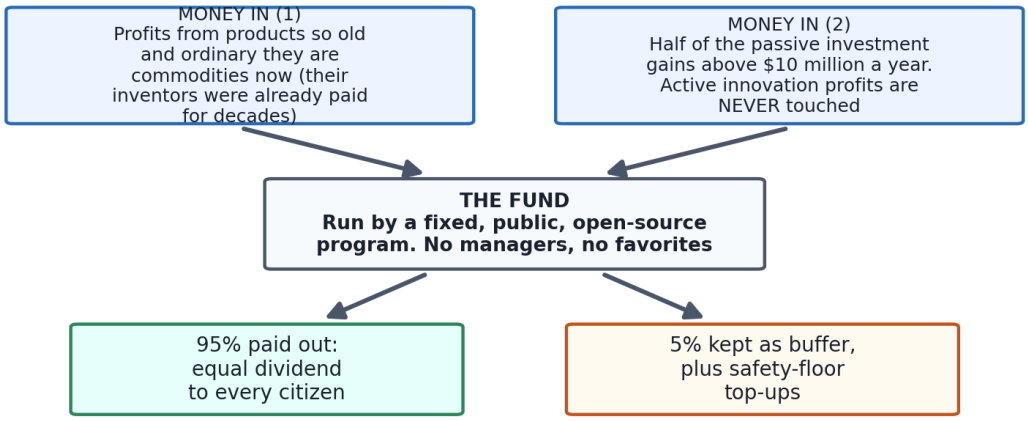

The citizen-owned fund

Where does the fund's money come from? Two places, and the choice of those two places is the moral center of the design:

1. **Profits from commoditized products.** When something new is invented, its inventor keeps the profits for decades (see Stream 3). Eventually every invention becomes ordinary, a commodity, like the light bulb or the pocket calculator. Profits from that mature, ordinary stage flow to everyone.
2. **A share of large passive investment returns.** Half of investment gains above 10 million dollars a year (rising for the very largest, oldest fortunes) go to the fund. Note what this is NOT: it is not a tax on building companies, inventing things, or working. It touches only money that money makes by sitting there.

The paper's phrase for this is worth quoting: the fund is **capitalism completed, not socialism introduced**. Nothing is confiscated (existing wealth is explicitly untouched; only its future passive returns above the threshold are shared). Everyone becomes a shareholder in the automated economy, playing by the same public rules. It levels the playing field, not the outcomes.

Stream 3: Innovation rewards

If you and your twin invent something people want, you keep outsized profits for **25 to 40 years**, patent-length protection. The paper's principle: reward the creators of new value for a long time; share only what has become passive or commonplace. As the paper puts it: wealth already compounds on its own, it needs no policy help. Creation does.

A clever detail: protection fades based on the product's actual profit margins, not on lawyers' arguments. When a product's margins compress to ordinary commodity levels, protection has done its job and expires on its own. There is no committee to lobby and no "new and improved version 2.0" trick to reset the clock, because a relabeled commodity earns commodity margins and loses protection automatically.

The floor

Beneath all three streams sits a guaranteed income floor, about 40,000 dollars per household in today's money and rising over time. If your three streams ever total less, the fund tops you up. In the simulations, top-ups are needed only during the transition years; by mid-century nobody needs them, because the three streams outgrow the floor on their own.

The floor is indexed to a realistic basket of what people actually buy, with weights that update automatically from real transaction data. That sounds like bureaucratic trivia; Section 6 explains why the paper's own testing proved it is not.

---

4. The math, without pain

The paper contains equations. Every one of them is an ordinary sentence wearing symbols. This section translates them all.

The economy's engine room: Y = K^a (A x L)^(1-a)

This is the **growth accounting** equation economists have used since Robert Solow invented modern growth theory in 1956. It just says: what an economy produces (**Y**) depends on its tools and machines (**K**, for capital), how good its technology is (**A**), and how many productive workers it has (**L**, for labor). The little exponent **alpha** says how much of the credit goes to machines versus workers.

> **Who found this?** Robert Solow won the Nobel Prize for this framework ("A Contribution to the Theory of Economic Growth," 1956). It remains the standard starting point for long-run growth arithmetic. The paper uses it in careful "index form," which avoids a units problem that tripped up an earlier draft, and openly brackets the biggest argument about it (what happens to alpha under automation) rather than betting everything on one side. Section 6 returns to that argument.

Using this arithmetic with deliberately cautious assumptions, the paper sketches a 2075 U.S. economy of 227 trillion dollars (about 9 times today), or 487 trillion under more optimistic technology (about 19.5 times). These envelope numbers are labeled scenario analyses, not predictions, and every one is checked by computer for internal consistency.

The money equation: M x V = P x Y

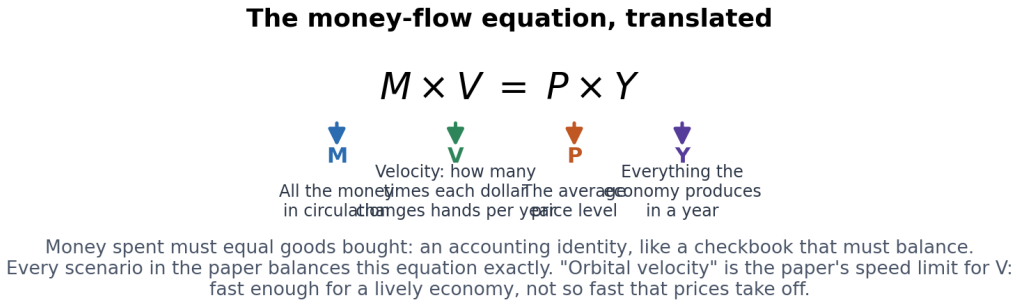

The money-flow equation

This one, from economist Irving Fisher (1911), is an accounting identity, meaning it must balance the way a checkbook must balance. Money in circulation (**M**) times how fast it changes hands (**V**, velocity) equals prices (**P**) times everything produced (**Y**). Every monetary scenario in the paper balances it exactly.

> **Who found this?** Irving Fisher formalized the "quantity theory of money" in *The Purchasing Power of Money* (1911). Milton Friedman and Anna Schwartz's monumental *A Monetary History of the United States* (1963) showed how money mismanagement deepened the Great Depression; Ben Bernanke's *Essays on the Great Depression* (2000) extended that lesson (he later ran the Federal Reserve during the 2008 crisis). McCandless and Weber (1995) documented that across countries, printing money faster mostly just raises prices. Robert Lucas (1972) and Thomas Sargent (1982), both Nobel laureates, showed how expectations

shape inflation. The paper cites this literature to show its monetary scenarios respect a century of hard-won lessons.

Orbital velocity: the speed limit

The paper adds one original concept here: **orbital velocity**. In a post-job economy money moves fast (instant AI-mediated payments), and fast-moving money can either energize an economy or overheat it into inflation. The paper defines a safe ceiling, like the speed a satellite needs to stay in orbit rather than crashing or flying off:

V_O = G_Y x I / S

In words: the sustainable speed limit for money (**V_O**) rises with real growth (**G_Y**) and confidence (**I**, a small bonus from supportive policy **R** and consumer confidence **C**), and falls when inflation runs above target (**S**, a stability brake that automatically lowers the ceiling when prices heat up). In the paper's scenarios, actual money velocity grows at about 2.2 percent a year, comfortably under ceilings of 4.7 to 6.4 percent: a lively economy in stable orbit, not a runaway.

The Gini number and demand breadth

You will see "Gini" throughout this guide. The **Gini coefficient** is the standard measure of inequality, running from 0 (everyone has exactly the same) to 1 (one person has everything). Today's United States income Gini is roughly in the 0.4 range; a Gini of 0.7 or 0.8 would be an extremely concentrated economy.

The paper defines **demand breadth** simply as B = 1 minus Gini. High breadth means spending power is widely spread; low breadth means it is bottled up at the top.

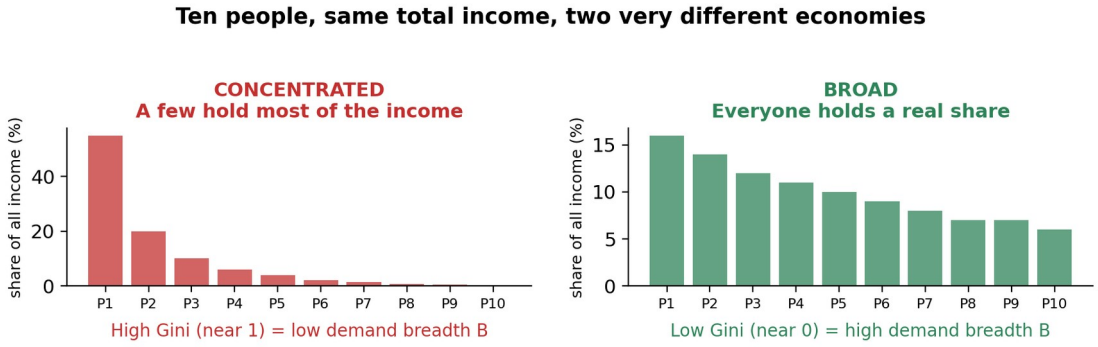

Same income, two different economies

The growth equation: the whole theory in one line

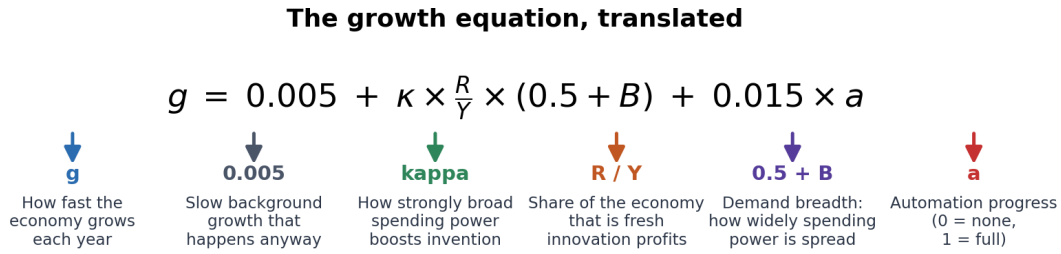

The growth equation, translated

Each year in the simulation, the economy grows by: a small baseline, PLUS the innovation term (the fraction of the economy that is fresh innovation profit, times demand breadth, times a strength dial called **kappa**), PLUS an automation bonus. The innovation term is Clutch Economics in mathematical form: invention drives growth, and invention pays off in proportion to how many people can buy.

Everything hinges on that dial, kappa. If kappa were zero, spreading money around wouldn't affect growth at all. The paper measured kappa from real-world data across 146 countries and got **0.37**, right in the middle of what prior research implies. It also honestly reports the measurement's limits. The companion guide to the peer-review paper covers that evidence in depth; the short version is in Section 6 below.

---

5. The computer model: a small nation in a bottle

You cannot run a 50-year experiment on a real country. So the paper builds an **agent-based model**: a simulated population of thousands of households (scaled up to represent 250 million American adults), each with its own wealth (matched to the real U.S. wealth distribution), its own inventive talent (a few people get a lot, most get a little, matching reality), and its own decisions, simulated year by year from 2025 to 2075. Every simulation is run many times with different random rolls, and the paper reports the middle outcome, so no result depends on a lucky run.

The model's crucial feature: **invention is a choice inside the model**, not an assumption. Simulated people launch ventures when the expected reward justifies it. That means if a policy kills the reward for inventing, the model will show invention collapsing. The design cannot cheat.

To keep itself honest, every experiment must pass the **Innovation Guard**: if any policy setting drops new-venture creation below 90 percent of what a completely free economy would produce, that setting FAILS, no matter how fair its outcomes look.

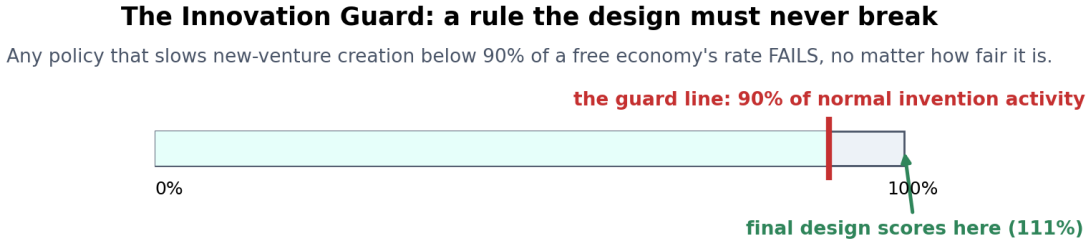

The Innovation Guard

The headline discovery

The model compares two worlds with identical technology and identical people. In one, the Clutch design (twins, dividend, innovation rewards, floor). In the other, hands-off laissez-faire: AI earnings simply flow to whoever owns the AI capital.

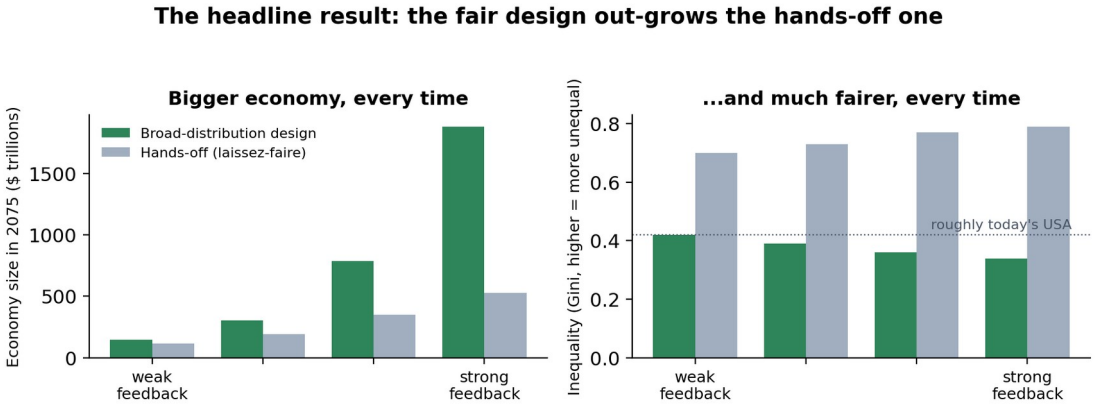

The headline result

The fair design produces a LARGER economy at every tested strength of the demand feedback, and the gap widens as the feedback strengthens. At the dial's midpoint the design's 2075 economy is roughly twice the size of the hands-off one. Meanwhile inequality stays near today's levels (Gini around 0.36 to 0.42) instead of soaring (0.70 to 0.79), and zero households sit below the floor.

Why does fairness out-grow concentration? Two reasons the hands-off world structurally lacks. **Bigger prizes:** when millions can buy your invention, inventing pays more. **More players:** when everyone has income and a twin, far more people can afford to attempt a venture at all. Concentration starves both.

And the pie picture matters as much as the pie chart:

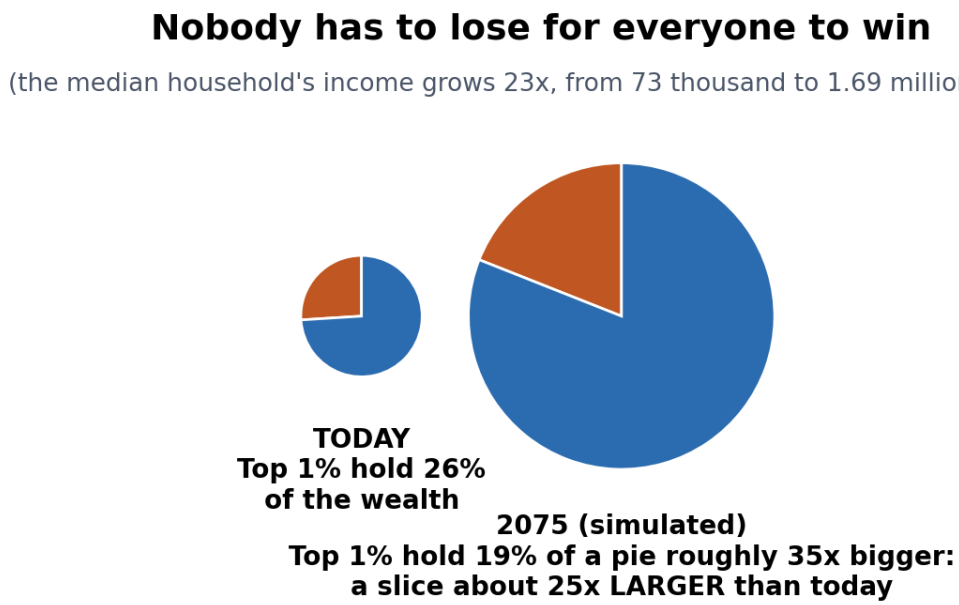

Nobody has to lose

The top 1 percent's SHARE shrinks from 26 percent to about 19 percent, but the pie grows roughly 35-fold, so their absolute wealth is about 25 times larger than today. The median household goes from 73 thousand to 1.69 million dollars a year. This is the paper's answer to zero-sum thinking: nobody needs to be torn down.

Three findings that surprised the author

Worth knowing because they show the model is a real test, not a rubber stamp:

1. **Long rewards won.** The author originally favored making innovation profits fade quickly (it sounds fairer). The model showed growth collapsing with fast decay while inequality barely improved, because fairness was already handled by the license and the fund. The design changed to patent-length rewards because the model said so.
2. **A prize add-on died.** An earlier draft let breakthrough inventors earn extended protection periods. The model showed the extension bought literally zero extra innovation while concentration rose. Deleted.
3. **The flattering number got demoted.** The paper's own real-income claims originally used a frozen 2025 shopping basket. A two-sector test (based on economist William Baumol's famous "cost disease": machines make goods cheap, but human services stay expensive) showed the honest number against an updating basket is 1.5-fold real growth for the median household, not 3.3-fold. The paper now reports BOTH numbers and redesigned the income floor around the honest one.

   Who found this? William Baumol ("Macroeconomics of Unbalanced Growth," 1967) explained why haircuts and violin lessons get relatively dearer forever even as televisions get cheap: you cannot automate the human touch without

losing what you were paying for. The paper embraces the implication instead of hiding it.

The stress tests

What if capital takes over? Most automation economists (notably Daron Acemoglu and Pascual Restrepo) expect machines' owners to capture a growing share of income. Skeptics would say that dooms a plan built on twin "labor" income. The model ran that exact world: the design got BETTER, not worse, because swollen corporate profits flow into the citizens' fund and the dividend automatically takes over from twin income. The design has a built-in shock absorber for precisely the future the mainstream predicts.

What if the transition is ugly? With automation arriving seven years early and twin adoption lagging, households below the floor peak at 7.3 percent in the early 2030s, then recover to zero by 2045. Uncomfortable, but nothing breaks, and the fund never goes broke.

Which assumptions actually matter? Wiggling every parameter up and down by half its value moves the results much less than the kappa dial does. Translation: there is one load-bearing number in the whole framework, the strength of the demand-breadth feedback, and that is exactly the number the paper went out and measured from data. Its best estimate: 0.37, with all results re-checked across the range 0.10 to 0.70.

---

6. The battle-testing: three rounds of paid sabotage

This is the part of the paper with, as far as we know, no precedent in economics writing. Instead of waiting for critics, the design hired its own.

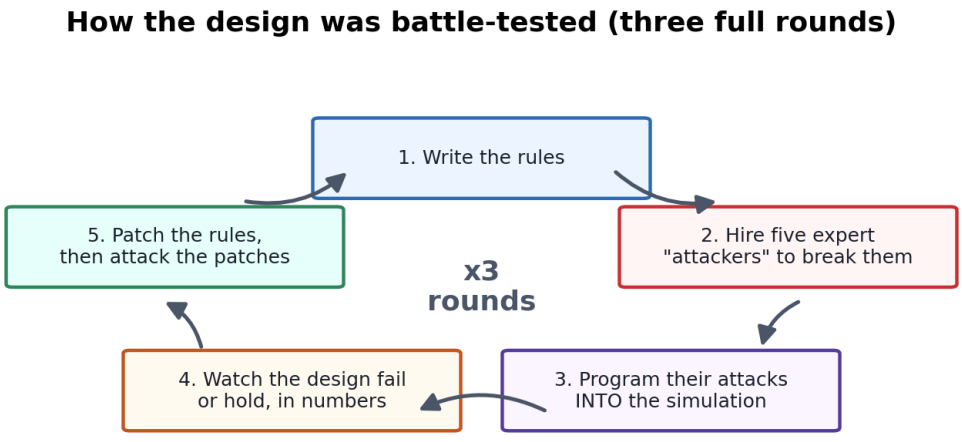

The red-team cycle

Five adversarial expert personas were convened three separate times: a regulatory-arbitrage attorney (a lawyer who finds loopholes for a living), a platform monopolist (thinks like a tech giant), a financial engineer (builds exotic contracts), a political lobbyist (captures institutions slowly), and a shadow economist (runs gray markets). Their job: break the design. Then, crucially, **every attack they invented was programmed into the simulation** and run, so the damage is measured in numbers, not debated in words. Only then were the rules patched, and the next round attacked the patches.

What they found, in plain words

Round one's killer: "The license can't be sold? Fine. I'll just buy the citizen's future INCOME forever, bundle it with an advisory contract that tells the twin what to do, and I own the twin in everything but name." Four of the five attackers independently found this. Simulated, it re-creates oligopoly within a decade: advisors capture 36 percent of twin income, the top 1 percent reach 45 percent of wealth, and the economy is less than half its patched size. The patch: income can only be pledged 3 years and 35 percent ahead, larger deals are legally void, and any service directing more than 0.4 percent of all twins gets structurally broken up. Result after patching: advisory capture 5 percent, top 1 percent at 17 percent.

Also round one: unlicensed offshore "gray" AI fleets undercutting citizens' twins (answered with a verification system: buyers of AI labor get tax safe harbor only for attested, licensed work, with privacy preserved and no central registry), and the rich splitting fortunes across family trusts to duck the 10-million-dollar threshold (answered with look-through rules that trace to the real owner).

Round two: attack the new rules. Finance found a way to rebuild corporate twin-fleets through revenue-based lending. The patch that survived three rounds of counterattack is the **economic-exposure rule**: anyone whose overall economic stake in a fleet's output exceeds 20 percent, through ANY instrument (shares, revenue deals, derivatives, leases, or even being the fleet's dominant customer), is simply taxed as an owner of that fleet. Not banned: taxed. Outside investment stays welcome (the model proved financing helps innovation); what dies is the untaxed disguise.

Round three: attack the exposure rule. The most damaging discovery in the whole project: an **offshore laundering** scheme that bought twin output cheap through unverified foreign middlemen. Simulated, it drained 35 percent of the twin pool offshore and was the only attack in all three rounds to break the Innovation Guard. The final patch: output sold to unverified foreign buyers pays an export withholding tax, refundable when the money verifiably comes home. Domestic business feels nothing; the laundry loses its margin.

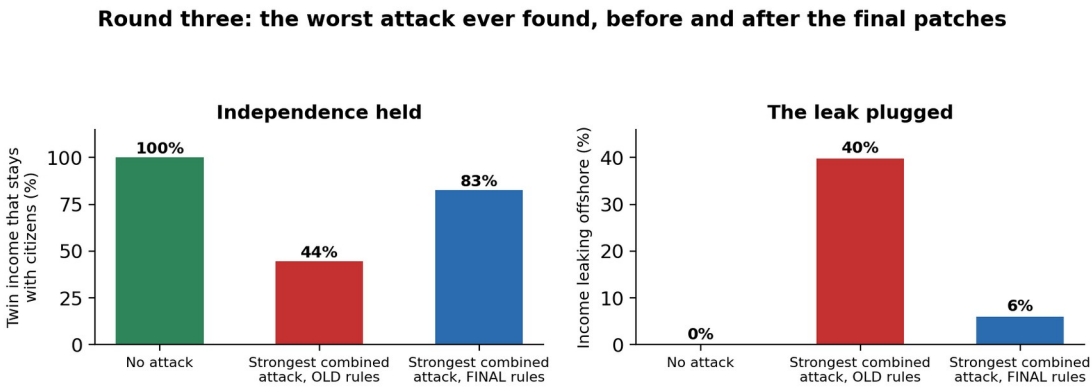

The worst attack, before and after

The punchline: the design got antifragile

Here is the result that exceeded the designer's own target. Run the finished design under the strongest COMBINED attack the three rounds could compose, and compare it to the same design with no attack at all. The attacked economy comes out with HIGHER output and a HIGHER citizen dividend (229 thousand versus 159 thousand dollars per person), because the exposure tax converts the attackers' capital into fund income, and their financing, pushed into legitimate channels, funds new ventures. The design does not merely survive its worst known attack. It profits from it. (The honest asterisks: independence of twin income dips to 83 percent, inequality ticks up two points, and the claim covers known attacks only. Someone may invent a new one. The paper says this plainly and publishes everything an attacker would need to try.)

---

7. The roadmap: no permission needed to start

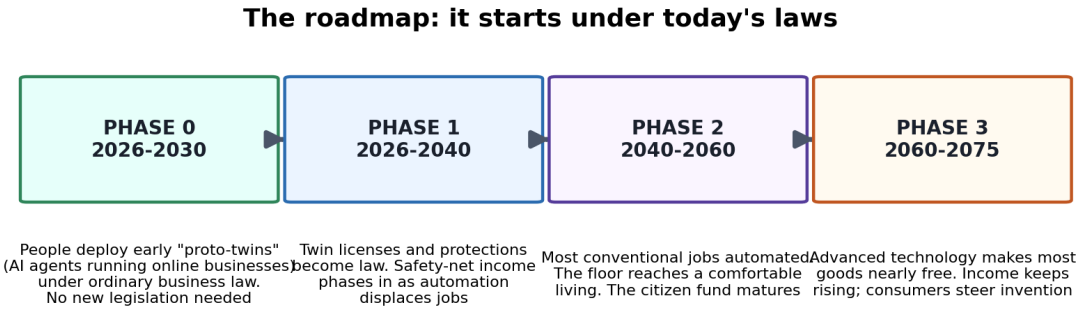

The roadmap

A crucial feature: **Phase 0 requires no new laws.** AI agents that run online businesses under ordinary business law already exist, today. People deploying "proto-twins" now generate the income data, the experience, and the constituency that make the later licensing debate winnable. The exotic technologies in the later phases (fusion power,

asteroid mining, molecular manufacturing) matter only for the ceiling of abundance, not for starting.

Real-world evidence anchors the near-term steps. Alaska has paid every resident an oil-fund dividend for four decades:

> **Who found this?** Damon Jones (University of Chicago) and Ioana Marinescu (University of Pennsylvania) studied the Alaska Permanent Fund dividend ("The Labor Market Impacts of Universal and Permanent Cash Transfers," 2022) and found the free money did NOT make Alaskans stop working. Norway's trillion-dollar sovereign fund shows disciplined rules-based wealth management at national scale. Michael Kremer (Nobel laureate), Jonathan Levin, and Christopher Snyder documented how prize-like "advance market commitments" (2020) got vaccines built for poor countries, the real-world precedent for the design's citizen-voted innovation prizes. And Andrew Yang's *The War on Normal People* (2018) brought automation-driven basic income into mainstream American debate; the paper differs from Yang by making earned twin income, not a government check, the primary channel.

---

8. Honest fine print

The paper closes with its own weaknesses, stated plainly. A few in ordinary language:

1. **The key number rests on cross-country comparisons.** Kappa = 0.37 comes from comparing countries against each other over decades. Tracking single countries over time cannot confirm it (the data is too sluggish), so a skeptic can defensibly believe kappa is near zero. If it were, the design would no longer out-GROW laissez-faire; it would merely tie on growth while being vastly fairer. The claim degrades gracefully; it does not flip.
2. **Everything is conditional on automation actually arriving** on something like the assumed schedule. The paper models the institutional choice, not the technology timeline.
3. **The attackers were thorough, not omniscient.** Fifteen expert passes plus an automated worst-case search found what they found. A genuinely new exploit class could exist. The code and attack reports are published as a standing invitation.
4. **Slow institutional capture** (lobbying that bends definitions over decades) is designed against (fixed formulas, randomized citizen panels, no standing committees) but cannot be fully simulated.
5. **It is a closed model of one country.** International coordination, migration, and trade are future work.

One more piece of honesty deserves its own paragraph: the paper's numbers are simulation outputs under explicit assumptions, verified for arithmetic consistency and open for anyone to re-run. They are labeled scenario analyses, never forecasts. Anyone who tells you they know 2075's GDP is selling something; the paper claims only "if these assumptions, then this arithmetic."

9. Who's who: every source in the paper, explained

In order of appearance by theme. Each entry: who they are, what the cited work says, and why the paper uses it.

On automation and its economics

- **Daron Acemoglu & Pascual Restrepo, "Robots and Jobs" (2020).** MIT and Boston University economists; the leading empirical study of what industrial robots actually did to U.S. local labor markets (reduced employment and wages where robots arrived). The paper cites it as ground truth that automation displacement is real, and engages its expectation that capital's share of income rises with automation.
- **Erik Brynjolfsson & Andrew McAfee, *The Second Machine Age* (2014).** MIT scholars; the standard book-length case that digital technology is transforming work as profoundly as steam and electricity did.
- **McKinsey Global Institute, "The Future of Work in America" (2023).** Consulting-firm research quantifying which U.S. jobs face automation, cited for displacement timelines.
- **Elon Musk, VivaTech remarks (2024).** Cited as a prominent industry voice publicly concluding an automated economy will need universal high income, evidence the conversation has reached decision-makers, not evidence in itself.

On abundance and advanced technology

- **Ray Kurzweil, *The Singularity Is Near* (2005).** The best-known argument that machine intelligence will surpass human intelligence within decades.
- **Peter Diamandis & Steven Kotler, *Abundance* (2012).** The case that technology repeatedly turns scarcity into plenty.
- **K. Eric Drexler, *Engines of Creation* (1986).** The founding book of molecular nanotechnology, the paper's Phase 3 manufacturing technology.
- **Federico Faggin, *Irreducible* (2024).** From the inventor of the first commercial microprocessor, an argument that consciousness is fundamental and not reducible to computation; cited where the paper discusses sentient AI citizens.

On growth theory and money

- **Robert Solow (1956).** Nobel laureate; the growth-accounting framework behind the paper's long-range arithmetic.
- **Joseph Schumpeter, *Capitalism, Socialism and Democracy* (1942).** The economist of "creative destruction": growth comes from innovators earning temporary outsized profits. The paper's innovation-rewards channel is deliberately Schumpeterian.
- **Irving Fisher (1911); Milton Friedman & Anna Schwartz (1963); Ben Bernanke (2000); Robert Lucas (1972); Thomas Sargent (1982); McCandless & Weber (1995).** The monetary canon (see Section 4): the paper's money scenarios are built to respect it.

- **Adam Smith, *The Wealth of Nations* (1776).** The founding text of market economics; cited for the "invisible hand," the principle that overall direction emerges from free individual choices, which the paper preserves deliberately: consumers steer, no planner does.
- **Martin Weitzman, *The Share Economy* (1984).** Harvard economist who proposed workers sharing in profits rather than fixed wages; an intellectual ancestor of universal ownership.
- **Thomas Sowell, *Basic Economics* (2014).** A widely read defense of market economics written for general readers; cited for the incentive fundamentals the design is careful never to violate.
- **William Baumol (1967).** The "cost disease" of human services (see Section 5); the reason the paper reports two real-income numbers and indexes the floor to a realistic basket.

On behavior and motivation

- **Daniel Kahneman & Amos Tversky, "Prospect Theory" (1979).** The most cited paper in behavioral economics: people feel losses about twice as strongly as gains. The paper uses it to explain zero-sum instincts, why voters reflexively believe someone must lose for others to win, and treats that bias as an education challenge rather than a truth.
- **Daniel H. Pink, *Drive* (2009).** Synthesizes psychology research showing people are motivated by mastery, autonomy, and purpose more than by paychecks; the paper's framework for what gives life structure after jobs (directing your twin, pursuing crafts, sophisticated consumption).

On real-world evidence for the mechanisms

- **Damon Jones & Ioana Marinescu (2022).** Alaska's universal dividend did not reduce work (see Section 7).
- **Michael Kremer, Jonathan Levin & Christopher Snyder (2020).** Advance market commitments: prize-like funding that demonstrably works (see Section 7).
- **Andrew Yang (2018).** Mainstreamed automation-driven income policy in U.S. politics.
- **World Bank, World Development Indicators.** The public database of country statistics used to measure kappa.
- **Emanuel F. Barros, companion materials (2026).** The paper's own published model code, red-team reports, specifications, and data, so anyone can re-run everything.

---

10. Small glossary of terms

- **Agent-based model (ABM):** a simulation built from thousands of individual decision-makers rather than one averaged equation; lets you test how rules play out among unequal, self-interested people.
- **Attestation:** privacy-preserving proof that a twin is unique, licensed, and its owner is alive, without any central registry watching everyone.

- **Commoditized:** an invention that has become ordinary, earning only routine profits (the light bulb today). In the design, that is when its profits start flowing to everyone.
- **Demand breadth (B):** 1 minus Gini; how widely spending power is spread.
- **Digital Twin:** your licensed personal AI agent earning income on your behalf.
- **Gini coefficient:** the standard 0-to-1 inequality measure (0 = perfectly even, 1 = one person has everything).
- **Innovation Guard:** the model's self-imposed rule: any policy that cuts new-venture creation below 90 percent of a free economy's rate automatically fails.
- **Kappa:** the measured dial (best estimate 0.37) for how strongly broad spending power boosts innovation-driven growth.
- **Laissez-faire:** "let it be": the hands-off benchmark where AI earnings flow to whoever owns AI capital.
- **Passive returns:** money earned by wealth sitting in investments, as opposed to money earned by creating something new.
- **Quasi-rent:** economist-speak for the temporary outsized profit an innovator earns before competition catches up.
- **Red team:** experts hired to attack your own system to find weaknesses before real adversaries do (a practice borrowed from security engineering).
- **Socialization (of returns):** routing a share of large passive gains to the citizen fund. In this design it never touches active innovation profits or existing savings.
- **Sortition:** selecting decision-makers by lottery (like jury duty) so positions cannot be bought or lobbied in advance.
- **UAI (Universal Automation Income):** the umbrella term for the combined income streams (twin income, dividend, innovation rewards, floor).
- **Velocity (V):** how many times a dollar changes hands per year.

This primer summarizes "Clutch Economics" v6 (July 2026) by Emanuel F. Barros. All numbers cited are from the paper's computationally verified simulations and scenario analyses; the model code and data are published with the paper. For the academic companion paper written for peer review, see the separate primer "The Peer-Review Paper, Explained for Everyone."