

~98%

of the value created by innovation goes to society, not to the innovators — the audited share kept by producers is 2.2 percent.

NORDHAUS · NBER W10433 · US NON-FARM ECONOMY 1948–2001

# The Value Engine

What Capitalism Pays Society — and How Agentic AI Multiplies It

## In Brief: The Ninety-Eight Percent Machine

Capitalism's winners keep about two cents of every dollar of value their innovations create. The other ninety-eight escape to society — as lower prices, better products, and goods that did not exist before. This report audits that payout across 106 primary documents, then asks what happens when agentic AI upgrades the machine that makes it.

**I**n 2004, William Nordhaus went through the accounts of the American non-farm economy from 1948 to 2001 and asked a simple question: of all the value created by innovation, how much did the innovators keep? The answer was **2.2 percent** — a central estimate in a range of 0.8 to 3.3 percent, with Schumpeterian profits amounting to just 3.8 percent of corporate profits. Everything else flowed to people who never signed a contract with the innovator.

That number inverts the standard story. The standard story says markets extract value from society and concentrate it at the top. The audited story says the market is a machine that **forces its winners to give away roughly 45 times their profit** in surplus (ENSI estimate:  $1 \div 0.022$ ). The mechanism is not benevolence but competition, imitation, and patent expiry: rivals copy, prices fall, and the surplus escapes to consumers.

The machine's total output is not small. The world produces roughly **\$150–180 trillion a year**, in purchasing-power terms, above what a pre-industrial world of eight billion people would produce (ENSI estimate from the Maddison Project series) — frontier income per head is up about 30-fold since 1800. The natural experiments answer the attribution question: Korea split in 1948 and the market half now out-produces the planned half by an order of magnitude; the freest quartile of economies produces 7.6 times the GDP per capita of the least free (\$52,877 vs \$6,968, Fraser 2024). Assign even half the modern flow to the institutional package and ~\$75–90 trillion a year exists because the rules allow prices, property, and entry.

And the receipts reached the bottom. Extreme poverty fell from 38 percent of humanity in 1990 to **8.4 percent in 2019** — more than a billion exits while population grew by 2.3 billion. China's market reforms account for roughly 800 million of them.

**2.2%**

Share of innovation's social surplus kept by producers, US 1948–2001

NBER w10433 • Nordhaus 2004

**\$150–****180T/yr**

Output above the pre-industrial baseline, PPP • ENSI estimate

Maddison Project • ENSI Calculations \$1

**7.6<sub>x</sub>**

GDP per capita, freest vs least-free quartile of economies

Fraser EFW 2024 • Ch. 1

**8.4%**

Extreme poverty in 2019, down from 38% of humanity in 1990

World Bank • PSP 2022

## The best parts are counted at zero

The audit has a measurement problem that biases every number downward. GDP records transactions; the surplus escapes uncounted. Choice experiments put the median user's price for giving up search engines at **\$17,530 a year**, email at \$8,414, digital maps at \$3,648 — services GDP values at roughly their advertising revenue (Brynjolfsson-Collis-Eggers, PNAS 2019). The uncounted US digital surplus plausibly runs \$5–8 trillion a year (ENSI estimate; willingness-to-accept runs above willingness-to-pay, and even halved the figure stays in the trillions). Fold in mortality, leisure, and inequality, and welfare-relevant growth has run roughly 1.5 times measured consumption growth in catching-up economies (Jones-Klenow). **Every headline number in this audit is therefore a floor.**

## The dominant term — and the honest debits

Rank the machine's seven mechanisms by value created and one towers. A dollar of R&D returns **\$13.30 of social benefit** at baseline parameters — \$5 under deliberately conservative assumptions, \$20+ counting health spillovers — a social rate of return near 67 percent against private returns of 10–15 percent (Jones-Summers, NBER w27863). Because private actors capture so little, the world persistently buys less innovation than the social return justifies: capitalism's greatest asset class is permanently on sale.

The debits are entered at full strength, then netted. The global top 10 percent takes about 52 percent of income; the China Shock cost 2.0–2.4 million American jobs with decade-long local scarring; average US markups drifted from 18 percent over cost in 1980 to 67 percent by 2014; the high-growth startup premium collapsed from 16 to 4 percent after 2000. These are real, bounded, and — this report will argue — **maintenance failures of competition, not refutations of it.** The credits run one to two orders of magnitude larger.

### What to know in 5 points

- 1 **The giveaway.** Innovators keep ~2.2% of the value they create; competition hands ~98% to society — a ~45x surplus multiplier on innovator profit (ENSI estimate).
- 2 **The counterfactual.** The market-institutional package produces \$150–180 trillion a year above the pre-industrial baseline (ENSI estimate, Maddison-based).
- 3 **The engine.** Innovation is the dominant mechanism: \$1 of R&D returns \$5–13.30 to society — a ~67% social rate of return (Jones-Summers).
- 4 **The multiplier.** Agents amplify all seven value mechanisms at once — a conditional dividend of \$10–20 trillion a year within a decade (ENSI estimate).
- 5 **The hinge.** Who keeps the new surplus is decided by market structure, not by the technology.

## 1B+

exits from extreme poverty,  
1990–2019 · World Bank

## +2.3B

population growth over the  
same period

## 2x

life expectancy since 1900,  
frontier economies · Preston

## \$5–8T

uncounted US digital surplus  
per year · ENSI estimate

## The multiplier and the bet

The audit is not the news; the news is what happens next. Every mechanism by which markets create value — prices, competition, innovation, entrepreneurship, capital allocation, trade, institutions — is at bottom an **information-processing or transaction-cost machine**. Agentic AI is a general-purpose collapse in the cost of exactly those two things. The macro debate is one disagreement wearing two spreadsheets: Goldman Sachs projects a ~7 percent global GDP uplift over ten years; Acemoglu bounds the effect at 0.71 percent of total factor productivity. Acemoglu's arithmetic assumes AI stays a tool applied to fixed tasks; Goldman's assumes it reorganises tasks, firms, and markets. The tool scenario yields a rounding error. The actor scenario yields a second engine.

The evidence floor already tilts one way. Support agents with an AI assistant resolved **14 percent more issues per hour — 34 percent more for novices** (5,179 agents, Brynjolfsson-Li-Raymond). Writing tasks ran 37 percent faster at measurably higher quality (Noy-Zhang). Eighty percent of US workers have at least a tenth of their tasks exposed; agent tooling lifts the share of quickly completable tasks from 15 percent to roughly half (Eloundou). Sum the mechanism-by-mechanism arithmetic and the conditional agentic dividend lands at **\$10–20 trillion a year within a decade** (ENSI estimate).

Conditional — because the Nordhaus question returns. The 2.2 percent producer share was never a law of nature; it was an achievement of antitrust, patent expiry, open science, and low entry barriers. The agentic era re-runs the experiment: will the producers of intelligence keep 2 percent, or 40? The warning light is already on — markups rose from 18 to 67 percent over cost before a single agent transacted. The to-do list is therefore institutional, not technological: agent-readable markets, open-agent access, identity and liability law, micropayment rails, telemetry before strategy.

“

*The market is a machine that forces its winners to give away roughly forty-five times their profit — to people who never signed a contract with them.*

THE VALUE ENGINE · ENSI · FROM NORDHAUS (NBER W10433)

**+7%**

Goldman: global GDP over 10 yrs

**≤0.71%**

Acemoglu: TFP bound over 10 yrs

**+14/34%**

field productivity, all/novice · NBER w31161

**\$10–20T**

conditional agentic dividend/yr · ENSI estimate


### THE SPINE OF THE REPORT

**"Capitalism is a machine that forces its winners to give away ninety-eight percent of the value they create — and agentic AI is the first technology that upgrades the machine itself; whether society keeps the new ninety-eight percent is decided by the rules we write in the next five years."**

# The Audit in Numbers

Twelve figures carry the whole argument. Three charts show the split, the poverty event, and the live bet. Every value traces to a library document; estimates are labelled.

FIG. 1 – THE SPLIT



| ANATOMY OF THE SPLIT                        | VALUE                  |
|---------------------------------------------|------------------------|
| Producer capture of innovation surplus      | 2.2%                   |
| Estimation range                            | 0.8–3.3%               |
| Schumpeterian share of corporate profits    | 3.8%                   |
| Implied giveaway multiplier · ENSI estimate | ~45x                   |
| Sample                                      | US non-farm, 1948–2001 |

**Reader note:** competition, imitation, and patent expiry pass 97.8% of innovation's surplus to consumers — the producer's sliver is the thin dark wedge. · NBER w10433 · Nordhaus 2004

2.2%

Innovation surplus kept by producers (range 0.8–3.3%)

NBER w10433 · Nordhaus 2004

~45<sub>x</sub>

Surplus handed to society per \$1 of innovator profit · ENSI estimate

ENSI Calculations §2 · 1÷0.022

\$150–180T

Annual PPP output above the pre-industrial baseline · ENSI estimate

Maddison Project · Calc. §1

~30<sub>x</sub>

Rise in frontier real income per head since 1800

Maddison · McCloskey

\$13.30

social benefit per \$1 of R&D, baseline · Jones-Summers

67%

social IRR of R&D vs 10–15% private · NBER w27863


\$14.5–38.6T

social value of global R&D/yr · ENSI estimate

\$5

per \$1 R&D even on conservative assumptions

FIG. 2 – THE POVERTY EVENT



**Reader note:** the share of humanity in extreme poverty (\$2.15/day, 2017 PPP) fell from 37.8% to 8.4% in one generation — the fastest poverty decline in recorded history. · World Bank, Poverty & Shared Prosperity 2022





FIG. 2B – WHAT GDP COUNTS AT ZERO



**Reader note:** median annual compensation users demand to lose each free service — all of it invisible to GDP. · PNAS 2019 · Brynjolfsson-Collis-Eggers

FIG. 3 – THE LIVE BET


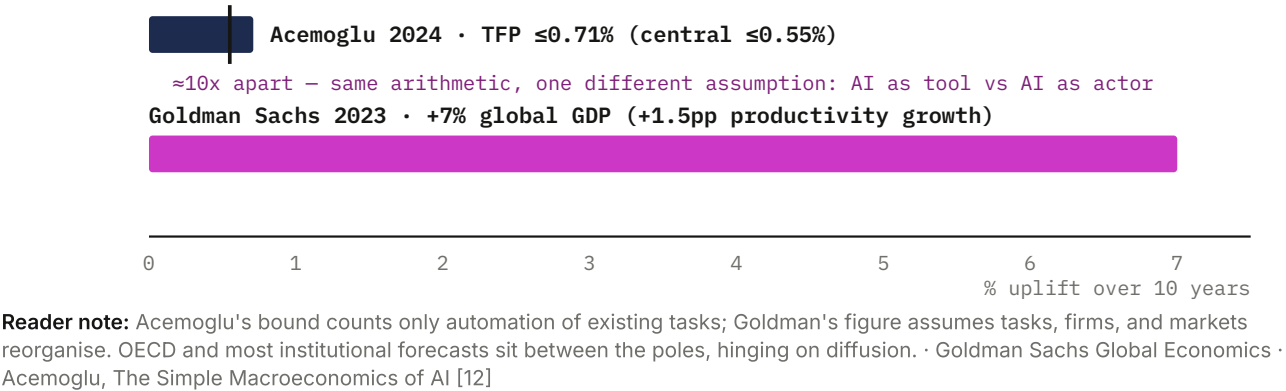


FIG. 3B – THE EVIDENCE FLOOR (ALREADY MEASURED)



**Reader note:** each bar is a separate study — resolved issues/hour (Brynjolfsson-Li-Raymond), task time at +0.45 SD quality (Noy-Zhang), and the share of tasks completable quickly once models get agent tooling (Eloundou). The floor is real before the bet settles.

Anthropic EI · ~36% of occupations use AI on ≥25% of tasks

57 / 43 · augmentation vs automation split

80% of US workers ≥10% task exposure · Eloundou

## THE VALUE ENGINE • 120 PAGES • FIVE MOVEMENTS

# Contents

## MOVEMENT A — ORIENTATION • 1-12

|                                         |    |
|-----------------------------------------|----|
| Cover .....                             | 1  |
| In Brief — Executive Summary .....      | 2  |
| The Dashboard .....                     | 5  |
| Contents .....                          | 8  |
| Method & Sources .....                  | 9  |
| The Reframe: The Donation Machine ..... | 10 |

## MOVEMENT B — THE AUDIT • 13-60

|                                       |    |
|---------------------------------------|----|
| Divider — Movement I: The Audit ..... | 13 |
| A1 The Counterfactual .....           | 14 |
| A2 The Split .....                    | 18 |
| A3 The Invisible Ledger .....         | 22 |
| A4 The Compounding Engine .....       | 26 |
| A5 The Poverty Event .....            | 30 |
| A6 The Welfare Multiplier .....       | 34 |
| A7 The Honest Debits .....            | 38 |
| A8 The Maintenance Thesis .....       | 42 |
| A9 The Distribution Question .....    | 45 |
| A10 The Systems Counterfactual .....  | 49 |
| A11 The Agentic Bet .....             | 53 |
| A12 Scenarios 2035 .....              | 57 |

## MOVEMENT C — THE MACHINE • 61-94

|                                             |    |
|---------------------------------------------|----|
| Divider — Movement II: The Machine .....    | 61 |
| The Machine Overview .....                  | 62 |
| R0 Institutions, the Operating System ..... | 65 |
| R1 Innovation & Creative Destruction .....  | 69 |
| R2 The Price System .....                   | 74 |
| R3 Competition .....                        | 78 |
| R4 Capital Allocation .....                 | 82 |
| R5 Entrepreneurship .....                   | 86 |
| R6 Trade .....                              | 90 |
| The Compounding Stack .....                 | 93 |

## MOVEMENT D — THE MULTIPLIER • 95-110

|                                                                 |     |
|-----------------------------------------------------------------|-----|
| Divider — Movement III: The Multiplier .....                    | 95  |
| The Amplification Thesis + Evidence Floor .....                 | 96  |
| Amplification Map I: Prices, Competition, Innovation .....      | 99  |
| Amplification Map II: Entry, Capital, Trade, Institutions ..... | 103 |
| The Agent-Native Economy .....                                  | 106 |
| The Nordhaus Question for AI .....                              | 109 |

## MOVEMENT E — PROOF &amp; ACTION • 111-120

|                                             |     |
|---------------------------------------------|-----|
| Divider — Movement IV: Proof & Action ..... | 111 |
| Country Receipts .....                      | 112 |
| The ENSI Operating Model .....              | 116 |
| The Evidence Base .....                     | 118 |
| Back Cover .....                            | 120 |

## How to read this report

**Depth 1** — cover, In Brief, and the Dashboard (pages 1-7): the whole argument in ten minutes. **Depth 2** — Movement B's twelve analyses: the audit in full. **Depth 3** — the mechanism catalog, the multiplier maps, and the country receipts. Magenta marks emphasis; **green** and **coral** appear only as positive and negative data. Back-of-envelope figures are always labelled "ENSI estimate."

# Method & Sources

Built on a downloaded, verified library: every number traces through one audited data pack to a primary document on disk.




The corpus spans the full perimeter: what capitalism is (Hayek, North), who captures the value (Nordhaus, Brynjolfsson), the growth and poverty record (Maddison, World Bank, UN), the seven mechanisms (NBER, Fraser, Kauffman, PIIE, CMA, CEA, OECD), the critiques (World Inequality Lab, Stiglitz, Ostrom), and the agentic frontier (Anthropic, DeepMind, BIS, x402, arXiv).

**How the numbers were audited.** Every figure lives in one data pack tagged to its source document; derived arithmetic sits in a calculations file with conservative and central variants. Verbatim quotations were admitted only after character-level verification against the library PDF; the unverifiable are paraphrased. Back-of-envelope aggregations are labelled "**ENSI estimate**"; source substitutions are disclosed on page 118.

| CONFIDENCE GRADE | MEANING                                                      | EXAMPLES                                                |
|------------------|--------------------------------------------------------------|---------------------------------------------------------|
| Very High        | Peer-reviewed and replicated, or official statistical series | Nordhaus 2.2% · World Bank poverty series · Maddison    |
| High             | Flagship institutional analysis, single authoritative source | Fraser EFW quartiles · PIIE trade payoff · Anthropic EI |
| Medium — BOE     | Single study, or labelled ENSI back-of-envelope aggregation  | \$150–180T counterfactual · \$10–20T dividend           |

FIG. 4 — SOURCE MIX (APPROXIMATE)



**Reader note:** counts are approximate; the full catalog appears on pages 118–119.

# The Donation Machine

Before the audit begins, one prior has to go: the belief that profit measures what capitalism takes from society. The audited accounts show the opposite — profit measures the roughly two percent of created value that the machine *lets* its winners keep.

The conventional view is coherent and wrong. It says: firms profit, therefore society pays; the richer the winners, the larger the extraction; GDP measures the benefit, so whatever GDP misses does not exist; and AI is simply the next rent-seeking industry arriving on schedule. Each link in that chain fails against a primary source. Profits are not the value — Nordhaus's accounting shows they are **about 2.2 percent of the value**, with the rest escaping to consumers through falling prices and better goods. Growth is not zero-sum — poverty collapsed fastest precisely where markets integrated, and more than a billion people exited extreme poverty while population grew. GDP is not the benefit — the median user values free search at \$17,530 a year and GDP records it near zero. Innovation does not mainly reward innovators — a dollar of R&D returns \$5–13.30 to society.

The reframe, then: capitalism is best understood not as an extraction machine with side benefits but as a **donation machine with a retention fee**. The donation is not chosen; it is forced — by imitation, entry, and patent expiry — which is exactly why it is reliable. Charity fluctuates with sentiment. The competitive giveaway has run for two centuries because no single winner can stop it.

| THE MYTH                                        | THE AUDITED REALITY                                                                                                                  | EVIDENCE                                 |
|-------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------|
| Profits measure what firms extract from society | Profits are ~2.2% of the value innovation creates; ~97.8% escapes to consumers as lower prices, better products, and new goods       | Nordhaus · NBER w10433 [02]              |
| The rich world got rich off the poor            | Growth is positive-sum: extreme poverty fell fastest where markets integrated — 38% → 8.4% of humanity, 1990–2019                    | World Bank PSP 2022 · Sala-i-Martin [03] |
| GDP measures the benefit                        | GDP counts free goods at zero; willingness-to-accept evidence shows trillions per year uncounted (\$17,530/yr for search alone)      | Brynjolfsson PNAS 2019 [02] · GDP-B [14] |
| Innovation mostly rewards innovators            | Social returns run 5–13x private returns; the social IRR of R&D is ~67% vs 10–15% private                                            | Jones-Summers · NBER w27863 [04]         |
| AI is just another software cycle               | Agents upgrade capitalism's own mechanisms — prices, competition, entry, allocation — not merely the products running on top of them | ENSI Reports 2–3 · [12][13]              |

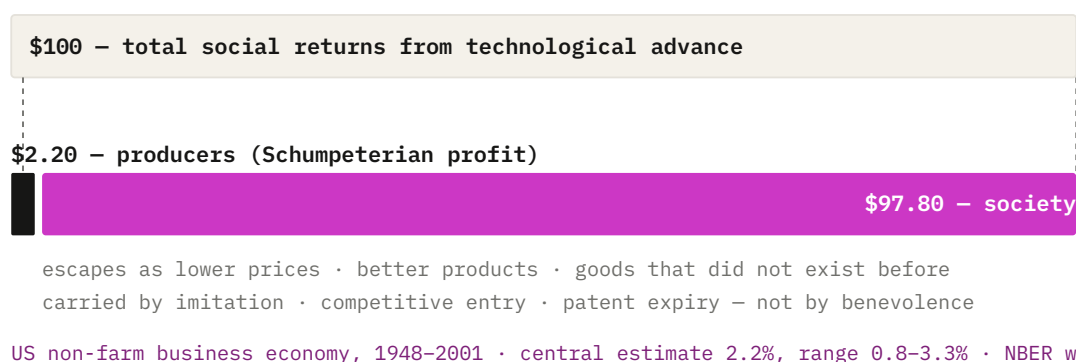
THE TEST THIS SECTION MUST PASS

If the myth column were true, the numbers in the next fifty pages could not exist. The audit is the falsification exercise; the reframe survives it.

## Anatomy of the 2.2 percent

Nordhaus's method deserves a page because the entire report leans on it. He defined **Schumpeterian profits** — profits arising from the appropriation of returns to innovation, as distinct from returns to routine capital — and estimated them for the US non-farm business economy across five decades of national accounts. Result: Schumpeterian profits averaged 3.8 percent of corporate profits, and producers captured 2.2 percent (range 0.8–3.3) of the total social surplus their innovations generated. The finding is conservative by construction: it uses measured productivity gains, and the new-goods literature (Hausman) shows measured gains undercount welfare from genuinely new products.

FIG. 5 – WHERE \$100 OF INNOVATION SURPLUS GOES



**Reader note:** the producer's cut is the thin dark bar. Everything to its right is the donation — surplus no contract assigned and no invoice recorded.

How does the surplus escape? Through four repeating steps, each documented later in the mechanism catalog:

- 1 **Rivals imitate.** Knowledge leaks; spillovers roughly double private returns to R&D (Bloom-Schankerman-Van Reenen), and the imitator's customers pocket the difference.
- 2 **Entrants price-compete.** New firms enter above-profit markets and bid margins down — the surplus migrates from the income statement to the price tag.
- 3 **Patents expire.** Temporary monopoly is the system's deliberate, time-boxed exception; at expiry the remaining surplus transfers to users permanently.
- 4 **New goods diffuse.** Products spread from early adopters to everyone; the welfare gain of novelty itself is the part GDP counts worst (Hausman).

“

*Only a miniscule fraction of the social returns from technological advances over the 1948–2001 period was captured by producers, indicating that most of the benefits of technological change are passed on to consumers rather than captured by producers.*

WILLIAM D. NORDHAUS · SCHUMPETERIAN PROFITS · NBER W10433, 2004 · VERIFIED VERBATIM

# The 45x giveaway — and why it is an achievement, not a law

THE GIVEAWAY MULTIPLIER • ENSI ESTIMATE

For every **\$1 billion of profit** innovators earned from genuine innovation, roughly **\$44 billion of surplus** flowed to everyone else —  $1 \div 0.022 \approx 45x$ . No philanthropy in history approaches this transfer rate, and it runs continuously.

To feel the scale, take one firm. Apple's cumulative net income is on the order of \$1 trillion. If smartphone-era profits sit anywhere near the Nordhaus split, the consumer surplus created by the smartphone category is measured in the **tens of trillions of dollars** — an ENSI back-of-envelope figure, but one independently consistent with the willingness-to-accept evidence on page 6, where a single free service is worth five figures a year to its median user. The most profitable company of the era is, on the audited arithmetic, one of history's largest net donors — involuntarily.

Involuntarily is the operative word, and it cuts both ways. Competition forces the donation: no firm can hold surplus that imitation, entry, and patent expiry are structurally designed to pry loose. But that means the 2.2 percent split is not physics. It is the output of **maintained institutions** — antitrust enforcement, patent terms that actually expire, open science, low entry barriers. Where the maintenance lapses, the split drifts: average US markups rose from 18 percent over cost in 1980 to 67 percent by 2014 (De Loecker-Eeckhout) — the machine still creating value, the giveaway share quietly eroding. The reframe therefore ends in an obligation, not a celebration: **the giveaway is the thing the rest of this report exists to defend** — into an era in which the producers of intelligence itself will test it again.

~45x

surplus per \$1 of innovator  
profit • ENSI estimate

~\$1T

Apple cumulative net income  
(worked example, BOE)

3.8%

Schumpeterian share of  
corporate profits • Nordhaus

18→67%

US markup drift 1980–2014 •  
the warning light

| WHERE THE ARGUMENT GOES NEXT  | THE QUESTION IT ANSWERS                                                          | PAGES   |
|-------------------------------|----------------------------------------------------------------------------------|---------|
| Movement I — The Audit        | How much value, through what channels, with what honest debits — twelve analyses | 13–60   |
| Movement II — The Machine     | The seven mechanisms that create the value, ranked and dissected identically     | 61–94   |
| Movement III — The Multiplier | How agents amplify every mechanism — and the new economy above human bandwidth   | 95–110  |
| Movement IV — Proof & Action  | Country receipts, the ENSI operating model, and the evidence base                | 111–120 |

NBER w10433 • Nordhaus 2004

NBER w23687 • De Loecker-Eeckhout

PNAS 2019 • Brynjolfsson et al.

ENSI Calculations §2 • giveaway multiplier

MOVEMENT I · THE AUDIT

# What the machine actually paid

Twelve analyses, one discipline: before any verdict on capitalism, read its accounts — the counterfactual, the split, the surplus GDP never records, the debits at full strength — and then ask what happens to every line when agentic AI arrives.

**01** The Counterfactual — what a world without the machine produces

**02** The Split — who keeps the value innovation creates

**03** The Invisible Ledger — the surplus GDP counts at zero

**04** The Compounding Engine — social returns to R&D

**05** The Poverty Event — the billion-person exit

**06** The Welfare Multiplier — what the money actually bought

**07** The Honest Debits — the red ink, at full strength

**08** The Maintenance Thesis — failures as engineering faults

**09** The Distribution Question — who got the Great Enrichment

**10** The Systems Counterfactual — what the alternatives paid

**11** The Agentic Bet — Goldman v Acemoglu, same arithmetic

**12** Scenarios 2035 — four futures for the split

## ANALYSIS 01 • THE COUNTERFACTUAL


## 01 The Counterfactual

The modern market economy produces on the order of **\$150–180 trillion a year** of output that a pre-industrial world would not — the largest positive number in the human record, and one almost nobody has ever seen written down.

**E**very audit begins with a baseline, and the baseline is the part everyone forgets. For nearly the whole span of human existence the material life of the median person did not improve. A peasant of 1750 lived at roughly the standard of a peasant of 1200, or of Roman Egypt: a few dollars a day, half of all children dead before adulthood, life expectancy stuck in the thirties. The Maddison Project's reconstruction — GDP per head for more than 160 countries, annual back to 1820 — draws the same flat line through every pre-modern society. Stagnation was not a failure state. It was the normal state.

Then, around 1800, in a corner of northwestern Europe, the line bent and kept bending. Real income per head in the frontier economies rose roughly **30-fold** — McCloskey's "3,000 percent" — a break with all prior history that compounding at a mere 2 percent a year, finally sustained instead of lost to war, plague or expropriation, was enough to manufacture. The human mind is not built to feel a thirty-fold change; we round it away and treat the modern standard of living as a fact of nature rather than the most improbable event in the economic record.

FIG. 1.1 – THE HOCKEY STICK



**Nothing, then everything.** World GDP per capita, 1990 international dollars. Growth per head was close to zero for a millennium, then rose ~14-fold in two centuries. Maddison Project 2020 • Bolt & van Zanden.

Maddison Project • 2020 update

McCloskey • The Great Enrichment

ENSI • CALCULATIONS §1

The arithmetic of the counterfactual

The question this analysis answers is brutally simple: how much output exists each year **because the machine exists**? Not how much profit, not how much growth over last year — how much of the world's annual product would vanish if productivity returned to its pre-industrial trend line. The answer requires only two numbers and a subtraction, and both numbers are among the best-attested in economics.

The first number is what the world produces now. World output runs at roughly **\$110 trillion nominal, or \$160–190 trillion at purchasing-power parity** — the PPP figure being the right one for a welfare comparison, since it prices what people can actually buy. The second number is what the same world would produce at pre-industrial productivity. In 1820, on the Maddison reconstruction, world GDP per head stood near \$1,102 in 1990 international dollars. A planet of 8.1 billion people producing at that rate generates on the order of **\$9 trillion PPP a year** — roughly Japan and Germany combined, spread across the whole of humanity.

Subtract, and the flow above the pre-industrial baseline comes to **≈ \$150–180 trillion PPP per year** (ENSI estimate, CALCULATIONS §1). That is the gross size of the modern growth machine's annual payout: not a stock of accumulated wealth but a recurring flow, re-earned every year, roughly seventeen to twenty times the entire counterfactual world economy.

Three honest caveats belong next to the number. First, the counterfactual is stylised: a never-industrialised world would not hold 8.1 billion people at \$1,102 a head. The estimate is an accounting device for scale, not a simulation. Second, PPP conversions across two centuries carry real error bars — hence the range. Third, and most important, **not all of the flow is "capitalism."** Science, states, geography and culture all sit inside it. The counterfactual sizes the prize; it does not, by itself, allocate credit.

Allocating credit is a causal question, and the rest of this analysis answers it with the strongest evidence social science has: places where history ran the comparison for us — the same people, split by a border, under different rules — and instruments that isolate institutions from everything they travel with. Both point the same way: the market-institutional package of property rights, prices, competition and openness is the **binding variable**. The freest quartile of economies produces 7.6 times the income per head of the least-free; the settler-mortality instrument puts a factor of about seven on institutions alone; the Korean peninsula produced an order-of-magnitude gap between two halves of one nation. Even a deliberately conservative reading — assigning only **half** the modern flow to that package — leaves a sum without precedent. The pages that follow assemble the case: first the natural experiments, then the cross-country gradient, then the attribution itself.

\$110T

world GDP, nominal · ≈\$160–190T PPP

~\$9T

world output at 1820 productivity, 8.1B people · ENSI estimate

\$150–180T

annual flow above pre-industrial baseline · ENSI estimate

~30x

frontier real income per head since 1800 · Maddison, McCloskey

## The experiments history ran

Cross-country correlations invite every objection in the book — rich countries can afford good institutions, geography confounds everything. The decisive evidence is different in kind: cases where **the same people, on the same land, at the same moment** were split between institutional regimes, and one clean instrument that isolates institutions from everything they travel with. Together they function as the control group the world never volunteered to be.

### The Korean peninsula, 1948–

One nation, one culture, one geography, partitioned into two economic operating systems. Within two generations the market South and the planned North diverged into an **order-of-magnitude income gap** — on the order of 10:1 to 20:1 — visible from orbit as a line where the lights stop. No confounder survives this comparison: everything is held constant except the rules.

### The two Germanys, 1949–1990

The same experiment, run in the middle of Europe on the continent's most industrialised population. Four decades of separation produced a several-fold living-standards gap between West and East — same-people, same-geography evidence in the 10:1-to-20:1 family — closed only after reunification imported the West's institutions wholesale.

### Botswana, 1966–2000s

Landlocked, tropical, resource-dependent — the exact profile that predicts stagnation. Instead: among the **fastest sustained per-capita growth in the world over ~35 years**. Acemoglu, Johnson and Robinson's verdict is institutional — secure property rights and constraints on the executive, not the diamonds; resource-rich neighbours with weaker institutions squandered comparable windfalls.

### The settler-mortality instrument

AJR's colonial-origins study isolates institutions with a variable no modern economy chooses: the mortality colonial settlers faced centuries ago, which shaped whether extractive or inclusive institutions were installed. The result: institutions alone account for roughly a **7-fold difference** in long-run income per head.

Read together, the cases control for culture, continent, resources and starting income — and the variable that tracks the outcome every time is the institutional package. This is as close to a randomised trial as macro-history allows, and it points one way: the flow computed on the previous page is not a coincidence of geography or a dividend of culture. It is **produced by rules**, and the rules can be identified.

AJR · MIT AER 2001

NBER w9124 · Botswana


Maddison [03] · divided-country series

ENSI library · angle 15, case studies

The freedom gradient, and the honest half

The natural experiments establish causality at the extremes; the cross-country gradient shows the dose-response in between. Sort the world's economies into quartiles by economic freedom — property rights, sound money, open trade, light regulation — and income per head climbs the ladder step by step, ending **7.6 times higher** at the top than at the bottom.

FIG. 1.2 – INCOME BY ECONOMIC-FREEDOM QUARTILE



Each step up the freedom ladder roughly doubles income. GDP per capita (US\$, PPP), quartile averages. Fraser Institute · Economic Freedom of the World 2024, Ch. 1.

The objection writes itself: gradients can be bought at the bottom's expense. The same dataset answers it. The **poorest tenth** of the population in the freest quartile earns **\$7,610** a year — more than the *average* citizen of the least-free quartile (\$6,968), and seven and a half times the \$1,010 the poorest tenth receives where freedom is lowest. The gradient does not trade the poor's welfare for the average; it lifts both.

|                                                             |                                               |                                     |                                                         |
|-------------------------------------------------------------|-----------------------------------------------|-------------------------------------|---------------------------------------------------------|
| <b>\$7,610</b>                                              | <b>\$1,010</b>                                | <b>7.6x</b>                         | <b>~7x</b>                                              |
| income of poorest decile, freest quartile · Fraser EFW 2024 | income of poorest decile, least-free quartile | freest vs least-free GDP per capita | long-run income effect of institutions alone · AJR 2001 |

Now the attribution, stated with deliberate conservatism. The experiments imply same-people gaps of 10:1 and up; the instrument puts a factor of seven on institutions alone; the gradient runs at 7.6x across 165 economies. Any of these read literally would attribute the great majority of the modern flow to the institutional package. This audit does not need the majority. **Assign it half.**

THE HONEST HALF-ATTRIBUTION – ENSI ESTIMATE

Even crediting the market-institutional package with only **half** of the \$150–180T annual flow above the pre-industrial baseline — less than every natural experiment implies — the machine's attributable output is **~\$75–90 trillion per year**, re-earned annually. Every remaining conclusion in this report survives on the half.


A payout of that size forces the next question, the one the whole politics of capitalism turns on: **who keeps it?** The answer — measured, not asserted — is Analysis 02.

## ANALYSIS 02 • THE SPLIT

## 02 The Split

For every dollar of value post-war innovation created, the innovating firms kept about **2.2 cents**. The other ninety-eight went to everyone else. The fortunes we argue about are the visible sliver of a dividend that is mostly invisible.

The intuitive model of capitalism — the one that powers most political argument — is that firms create value and capture it, leaving the public whatever they are compelled to leave. The data invert this. Working through US non-farm business data across the entire 1948–2001 period, William Nordhaus asked what fraction of the total social returns to innovation the innovators themselves actually kept as profit, rather than losing to lower prices and better products for consumers. His answer: **2.2 percent**, ranging from about 0.8 to 3.3 percent across specifications. The public is not the residual claimant on the value markets create. **The public is the primary claimant, and the firm is the residual.**

**0.8–3.3%**

full range of the producer share across Nordhaus's specifications — the qualitative result survives every variant

NBER w10433 • Nordhaus 2004

**3.8%**

share of total corporate profits that were genuinely Schumpeterian — rents from innovation rather than from position

NBER w10433 • Nordhaus 2004

Nothing in the chart is a moral claim; it is an accounting identity about where the surplus landed. The interesting question is *why* a system built on the pursuit of profit pays out ninety-eight percent of the winnings to people who never signed the cheque — and whether that ratio is a law of nature or an achievement that can be lost. That is where this analysis goes.

## Why the machine forces the donation

The 2.2 percent is not benevolence, and it is not redistribution after the fact. It is what remains of a monopoly after competition has finished with it. Every genuine innovation begins life as a temporary monopoly — the innovator alone can make the thing, and prices it accordingly, capturing a large share of the value while exclusivity lasts. The system then dismantles that position through three mechanisms, each unglamorous, each relentless.

- 1 **Imitation.** Rivals reverse-engineer, copy and improve. Every imitator bids the price down toward cost, and every increment of the price fall is surplus transferred from the innovator's margin to the consumer's pocket — without a single law being passed.
- 2 **Patent expiry.** Even where the state grants exclusivity, it grants it on a timer. When the patent lapses, generics and clones arrive, and the remaining stream of value — usually decades of a successful product's life — flows at competitive prices to the public.
- 3 **Competitive entry.** Excess returns are a signal, and entrepreneurs answer signals. New entrants with better versions and lower costs compress margins economy-wide, so even un-imitated innovators must price as if the imitators were already here.

Run those three continuously over half a century and the integral is the Nordhaus number: the innovator's take shrinks toward a sliver, and the sliver is what we see. Nordhaus's companion estimate makes the point from the other side: genuinely Schumpeterian profits — returns attributable to innovation rents rather than to market position — amounted to only about **3.8 percent of corporate profits** over the period. Even inside the profit column, innovation rents are the minority entry.

Notice what the finding does to the standard political argument. The critique of capitalism prices the system by its visible winnings — the fortunes, the share prices, the quarterly

earnings — and asks whether society got a fair deal against them. But the visible winnings are the **two percent**. Arguing about capitalism by staring at its billionaires is auditing a company by reading only the CEO's expense account: not wrong as far as it goes, but missing ninety-eight percent of the ledger.

Notice also what the finding is *not*. It is not a claim that every fortune is earned — 96.2 percent of corporate profit, on the same estimate, is something other than innovation rent, and some of that is position, capture and the market-power drift audited in Analysis 07. Nor is the 2.2 percent a constant of physics. It is an equilibrium outcome of specific conditions: antitrust that keeps entry open, patents that actually expire, science that publishes, capital that funds attackers. Weaken those and the producer share climbs — a possibility that becomes the central question of this report's final movement, because agentic AI will re-run this equilibrium from scratch.

“

*Only a miniscule fraction of the social returns from technological advances over the 1948–2001 period was captured by producers ... most of the benefits of technological change are passed on to consumers rather than captured by producers.*

WILLIAM NORDHAUS · SCHUMPETERIAN  
PROFITS IN THE AMERICAN ECONOMY · NBER  
W10433 · 2004

The method behind the number matters, and the next spread prices its implications: what a 2.2 percent producer share means in dollars, worked through the most valuable company of the smartphone era — and why the measured split is still an *understatement* of society's take.

## Pricing the giveaway

Ratios anaesthetise; dollars wake the reader up. Invert the Nordhaus share and the split becomes a multiplier: if producers keep 2.2 percent, then every dollar of genuine innovation profit you can see implies roughly **forty-five dollars** of surplus you cannot — dispersed into lower prices, better products, saved time and new capabilities, collected by consumers who never sent the innovator a cheque.

### THE 45X GIVEAWAY MULTIPLIER – ENSI ESTIMATE

$1 \div 0.022 \approx 45$ . For every **\$1 billion** of profit earned from genuine innovation, roughly **\$44 billion** of surplus flowed to everyone else. Read any innovation fortune in the news through this lens: multiply by ~45 to see the invisible side of the transaction.

Apply it to the most scrutinised fortune of the era. Apple's smartphone-decade profits are the standard exhibit for the prosecution — the margin on the device in a billion pockets. The Nordhaus lens flips the exhibit: if profits of that scale sit anywhere near the historical split, they are the *small* number in the transaction.

### Worked example — the smartphone split ENSI estimate · illustrative

Apple cumulative net income, smartphone era  $\approx$  \$1.0 trillion

Nordhaus split: profit  $\approx$  2.2% of social surplus  $\rightarrow$   $\$1.0T \div 0.022$

Implied consumer surplus of the category  $\approx$  **tens of trillions of dollars**

Read as an order of magnitude, not a point estimate: the 2.2% is an economy-wide average applied here to a single firm, and some of Apple's profit is position rather than innovation rent. But even at a quarter of the implied figure, the surplus delivered to smartphone users dwarfs the fortune extracted from them — a conclusion the willingness-to-accept evidence in Analysis 03 reaches by an entirely independent route.

This is the arithmetic that the political argument about "big tech profits" never runs. It does not acquit any firm of any particular abuse — market power, privacy, lock-in are audited elsewhere in this report at full strength. It establishes something prior: the **direction of net flow**. The transaction between the innovating firm and society is, by two orders of magnitude, a transfer *to* society.

## ~45x

surplus to others per \$1 of  
innovation profit · ENSI  
estimate

## 2.2%

producer share of innovation  
surplus · Nordhaus 2004

## 3.8%

of corporate profits genuinely  
Schumpeterian · Nordhaus  
2004

## \$1T→\$45T

Apple worked example, order  
of magnitude · ENSI estimate

## The method, and why 2.2 percent is a ceiling on the producer share

A number this consequential must earn its keep, so state the method plainly. Nordhaus works from the production side: he measures the profits attributable to innovation — Schumpeterian rents, isolated from returns to capital and market position — and compares them against the total social value the innovations created, tracked through productivity and prices over the full 1948–2001 record. The producer share is the ratio. Built from fifty years of national accounts rather than case studies of famous inventions, it is the average over everything, failures and dead ends included — not a selection of flattering anecdotes.

The robustness checks matter as much as the point estimate. Across Nordhaus's specifications the share moves between roughly **0.8 and 3.3 percent**. What never moves is the qualitative result: under every variant, producers keep a few cents on the dollar and society keeps the rest.

There is, however, a deeper reason to treat 2.2 percent as a *ceiling* on the true producer share, and it comes from the measurement literature. The denominator — the total social value created — is computed from measured prices and output. But measured prices systematically miss the largest welfare event in a product's life: **its arrival**. Jerry Hausman's canonical NBER work on the valuation of new goods makes the point: the year before a genuinely new good exists, its effective price is

infinite, and the fall from "unavailable at any price" to "on the shelf" is a consumer-surplus gain that standard price indices never record.

Correct for that — as Hausman does for exemplar products — and true consumer welfare growth rises materially above the measured figure. Every such correction enlarges the denominator of the Nordhaus ratio while leaving the numerator, recorded profits, untouched. The arithmetic is one-directional: **the better we measure, the smaller the producer share gets**. The 2.2 percent is what the split looks like through an instrument that is blind to new-goods surplus; the true share is lower.

What has been established, then. First: the modern economy's headline payout, computed in Analysis 01, is not hoarded by the firms that generate it — competition, imitation and patent expiry force roughly ninety-eight percent of innovation's value out into the public's hands. Second: this outcome is a **mechanism, not a mercy** — it holds because specific conditions (open entry, expiring patents, published science, contestable markets) keep holding. Third: even the ninety-eight percent is understated, because our measuring instrument records the surplus of new and free goods at approximately zero.

That third point is not a footnote. It is a ledger of its own — trillions of dollars a year of value delivered to households and counted by GDP at nothing. Weighing it is Analysis 03.

### THE ONE-WAY CORRECTION

Every improvement in welfare measurement discovered so far — new-goods pricing (Hausman), free-goods valuation (Analysis 03), quality adjustment — moves value *from* the unmeasured column *into* society's take. No known correction moves value toward producers.

NBER w10433 · Nordhaus 2004

NBER · Hausman, Valuation of New Goods

ENSI · CALCULATIONS §2


## ANALYSIS 03 • THE INVISIBLE LEDGER

## 03 The Invisible Ledger

The median American would need to be paid **\$17,530 a year** to give up search engines — value the national accounts record at roughly **zero**. The largest line items in capitalism's payout are the ones GDP cannot see.

GDP counts what is bought. It is close to silent on what is gained. When a free search engine replaces a paid encyclopaedia, GDP falls; when a mapping app worth thousands a year to a household costs it nothing, GDP records nothing. The error runs in one direction — against the system being audited — so an honest audit must re-instrument. The cleanest instrument is the one Brynjolfsson, Collis and Eggers built in massive incentive-compatible online choice experiments: stop asking what people *pay* for a good, and ask what they would have to be **paid to give it up**.

FIG. 3.1 — WHAT GDP COUNTS AT ZERO



**The price of giving it up.** Median annual willingness-to-accept (WTA) to forgo each free digital category, US users. Brynjolfsson, Collis & Eggers • PNAS 2019.

Two readings of the chart, both correct. The comfortable one: free digital goods are a gigantic unrecorded gift. The uncomfortable one: our headline measure of prosperity has drifted so far from experienced welfare that its single most-used category — the search bar — is valued at seventeen thousand dollars a year by its median user and at approximately nothing by the statistical system that judges the economy delivering it.

PNAS 2019 • Brynjolfsson-Collis-Eggers

NBER w25695 • GDP-B

Stanford/AER • Goolsbee-Klenow

## Triangulating the invisible: four independent methods, one direction

Stated-preference experiments carry a known weakness — people answering surveys spend imaginary money — so the finding stands or falls on whether independent methods, built on different data and different assumptions, land in the same territory. Four do.

**Method one: revealed time.** Goolsbee and Klenow refuse to ask anyone anything. People demonstrate what the internet is worth by what they give up for it — hours of leisure that could have gone elsewhere. Pricing online time at the user's opportunity cost of leisure, they put the consumer surplus of internet access at roughly **2–3 percent of full income** — far above the sliver of measured spending it accounts for. No hypothetical dollars anywhere in the calculation; the valuation is carried entirely by observed behaviour.

### Method two: experimental GDP-B.

Brynjolfsson's GDP-B programme is the concrete proposal to add the benefits of free goods back into the ledger — the "B" is for benefits — using incentivised experiments where real users accept real payments to actually deactivate services. The flagship result: the median user demands about **\$42 a month** to give up Facebook, and counting that single platform's surplus would have added **0.05–0.11 percentage points per year** to measured welfare growth. One free app, properly counted, shifts the growth rate of the world's largest economy at the second decimal — and the app is not even in the top four categories of the WTA chart.

**Method three: quality-adjusted prices.** Byrne and Corrado, working inside the Federal

Reserve's measurement tradition, rebuild price indices for consumer digital services with honest quality adjustment — capacity, speed, capability per dollar. Measured this way, the effective price of digital services has collapsed for decades faster than official deflators record, which means real consumption growth has been correspondingly understated. The same invisible surplus, found this time in the price statistics themselves.

### Method four: the accounting framework.

Diane Coyle's ESCoE work formalises why the gap exists at all: the production boundary of the national accounts was drawn for a priced, industrial economy, so zero-price goods generate real household welfare that falls outside it by construction. This is not a fringe complaint against GDP; it is the settled diagnosis of the people who maintain it — the Stiglitz-Sen-Fitoussi line that statistics should measure the welfare of households, not only the volume of transactions.

Four methods, four data sources, one direction. The magnitudes differ — stated WTA runs highest, time-use and quality adjustment lower — but every instrument that has ever been pointed at the free-digital economy finds **large positive surplus where GDP records approximately nothing**. In an audit, when four independent estimators disagree on size but agree on sign and order, the honest entry is a range, not a refusal.

The next page books that range: what the invisible ledger plausibly sums to for one economy, with the caveats priced in rather than hidden.

### WHY THE DIRECTION SURVIVES EVERY CRITIQUE

Each method's known bias points a different way — surveys overstate, time-use depends on the leisure price, deflators lag quality — yet all four land on the same sign. The conclusion that survives: the measured economy understates what households actually receive.

## Booking the entry: what the uncounted surplus sums to

Aggregate the instrument readings for the United States alone, and label the arithmetic for what it is. Search's median WTA of \$17,530, carried across roughly **210 million adult users**, implies on the order of **\$3.7 trillion a year** of surplus from one product category. Extend across the full digital bundle — email, maps, streaming, social, the rest of the free stack — and the plausible total lands at **\$5–8 trillion a year** of consumer surplus that the national accounts record at approximately zero (ENSI estimate, CALCULATIONS §3). Set that against the measured digital-sector contribution to GDP, near \$2 trillion, and the invisible entry is two-and-a-half to four times the visible one.





FIG. 3.2 – VISIBLE VS INVISIBLE



**Most of the digital dividend is off the books.** US, annual flow. ENSI estimate from PNAS 2019 WTA figures; central bundle value shown at range.

### THE HONEST DISCOUNT – WTA RUNS ABOVE WTP

Willingness-to-accept exceeds willingness-to-pay, sometimes severalfold; stated magnitudes are orders of magnitude, not invoices. So apply a brutal haircut: **halve every figure**. The invisible ledger still books trillions per year, still exceeds the sector's measured GDP, and still points the same way.

Nothing here is new behaviour by the economy — it is old behaviour by the instrument. New goods have always entered the accounts at zero surplus; the free digital stack merely industrialised the blind spot. Which is why the correct reading of this entry is not "add \$5–8 trillion to the score" but something stronger and more disciplined: **every measured number in this audit is biased downward**.

## The floor principle

Three analyses in, a pattern has emerged that deserves to be named, because it governs how every remaining number in this report should be read. The counterfactual flow of Analysis 01 is computed from *measured* output. The 97.8 percent society-share of Analysis 02 is computed against *measured* surplus. And Analysis 03 has now shown that measurement itself is structurally blind to whole categories of value — free goods, new goods, quality — and blind in one direction only.



*Every number in this report is a floor.*

ENSI · THE VALUE ENGINE · THE AUDIT'S READING RULE

The rule has teeth. It means the \$150–180 trillion counterfactual understates the machine's true annual payout; on the corrections surveyed here, the unmeasured surplus plausibly adds another 50–100 percent on top of the measured flow (ENSI estimate, CALCULATIONS §10). It means the producer share of innovation value is at most 2.2 percent, and the giveaway multiplier at least 45. And it means the burden of proof in the political argument sits where few place it: the case *against* the machine must be made net of trillions in payouts its critics' own statistics never counted.

- ▶ **Free goods** — priced at zero, valued at thousands per user per year; \$5–8T/yr uncounted in the US alone (ENSI estimate; Brynjolfsson-Collis-Eggers, PNAS 2019).
- ▶ **New goods** — the fall from "unavailable" to "purchasable" is the largest welfare event in a product's life, and price indices never record it (Hausman, NBER).
- ▶ **Health & longevity** — the doubling of life expectancy appears in no output statistic; Analysis 06 prices it, and it may be the largest entry of all (Jones-Klenow, AER 2016).

The floor principle also disciplines this report's own side of the argument. Where ENSI aggregates run above what conservative methods support — WTA above WTP, stylised counterfactuals — they are labelled, haircut, and quoted as ranges. An audit that inflates its credits forfeits its debits; ours must survive the halving of every soft number, and does.

What the floor principle cannot tell us is *why* the machine keeps generating value to undercount — why the surplus is created at all, year after year, rather than exhausted. The answer is the machine's core reactor: ideas, and the extraordinary arithmetic of what a dollar of research returns to society. That is **Analysis 04 — The Compounding Engine**.

---

**2.2%**

producer share — the split,  
measured · Nordhaus

**\$17,530**

median WTA for search, per  
year · PNAS 2019

**2–3%**

of full income: internet  
surplus, time-use method ·  
Goolsbee-Klenow

**+50–100%**

plausible unmeasured value  
atop measured flow · ENSI  
estimate

---

## MOVEMENT I · THE AUDIT — ANALYSIS 04

## 04 The Compounding Engine

Every dollar the market system invests in research returns roughly **\$13.30 to society** — a 67 percent social rate of return no private portfolio has ever matched. The investor keeps a sliver; everyone else gets the compounding — and the engine still runs chronically under-fuelled.

### \$13.30

/ \$1

Social benefit per dollar of R&D, baseline economy-wide estimate

Jones & Summers ·  
NBER w27863

### 67%

Social internal rate of return on R&D — vs 10–15% private returns

Jones & Summers ·  
NBER w27863

### 2<sub>x</sub>

Social over private returns from technology spillovers alone

Bloom-Schankerman-Van Reenen · NBER

### \$2.9<sub>T/yr</sub>

Global R&D spend, 2025 — the engine's entire fuel line

CALCULATIONS §4

**T**he first three analyses counted value already delivered. This analysis audits **the machine that manufactures the next round**. Its headline is the most consequential ratio in the modern growth literature: one dollar invested in research and development returns, on the best economy-wide estimate, \$13.30 of social benefit in present value — about \$5 under deliberately conservative assumptions, more than \$20 once health gains and international spillovers are counted.

Expressed as an interest rate, the baseline is a **67 percent social internal rate of return**, against the 10–15 percent private investors earn on research-intensive assets. That gap is the defining fact about how a capitalist economy treats its most productive activity — and it is the phenomenon itself, not a flaw in the study design: firm-level work by Bloom, Schankerman and Van Reenen finds **social returns to corporate R&D run roughly twice private returns** from technology spillovers alone. Knowledge leaks — by imitation, employee mobility, patent expiry — and almost none of the return can be invoiced by the party who paid for the research.

**NBER w27863** Jones & Summers · Social Returns to Innovation

**NBER** Bloom-Schankerman-Van Reenen · Technology Spillovers

**Stanford** Bloom-Jones-Van Reenen-Webb · Ideas Harder to Find

## Five, thirteen, twenty: a return no market can arbitrage

FIG. 4.1 – SOCIAL BENEFIT PER \$1 OF R&D INVESTMENT, PRESENT VALUE



**Even the floor pays five to one.** The conservative case survives a 20-year diffusion delay and the assumption that only part of growth traces to measured R&D. Jones & Summers, NBER w27863.

The Jones–Summers figure is not a survey average of project-level studies; it is an economy-wide accounting exercise. If long-run growth in income per head comes from innovation, then the increment of GDP attributable to that growth *is* the return on the economy's total innovative effort — an accounting that nets out every spillover margin at once: imitation, business stealing, intertemporal displacement. The denominator is everything the United States spends on R&D across business, government, universities and non-profits: about **2.7 percent of GDP**. The entire innovative effort of the largest economy on earth is a rounding error against its own output, earning a return no pension fund will ever see.

Stress-test it and the number bends but does not break. Impose a 20-year diffusion delay and the benefit is still **\$4.90 per dollar** — the internal rate of return falls to 11 percent, still above any hurdle rate a market applies to anything. Count health gains, inflation mismeasurement and spillovers to the rest of the world, and the authors' own arithmetic runs to \$10, even \$20 per dollar — for the US economy alone.

The firm-level evidence agrees from below. Across US panel data for 1980–2001, Bloom, Schankerman and Van Reenen used firms' distinct positions in technology space and product space to separate knowledge spillovers from business stealing: both operate, **the spillovers dominate**, and the economy under-invests in R&D as a result. Nor is the return confined to private laboratories — the Dallas Fed finds high long-run productivity returns to *government* R&D as well (Fieldhouse & Mertens, 2023).

“

*Only a miniscule fraction of the social returns from technological advances over the 1948–2001 period was captured by producers, indicating that most of the benefits of technological change are passed on to consumers rather than captured by producers.*

**WILLIAM D. NORDHAUS** • SCHUMPETERIAN PROFITS IN THE AMERICAN ECONOMY • NBER W10433 • VERIFIED AGAINST SOURCE

## The under-fuelled engine

State the argument at its plainest. A rational investor funds research up to the point where the *private* return meets the hurdle rate. But the private return is a sliver of the social return: Nordhaus's audit of the entire post-war US record found innovators capturing about **2.2 percent** of the social surplus their innovations created (Analysis 2 installed that number as this report's lens). An investor who keeps two cents on the social dollar stops investing long before society would want them to. The consequence is arithmetic, not ideology: **the economy's highest-return activity is systematically under-supplied**.

The 67 percent social rate of return is itself the proof. In any private market, a return that size would be competed away within years — capital would flood in until the return fell to the hurdle rate. The fact that the social return on research has *persisted* at these levels for decades tells you no one can capture it. It is the standing arbitrage that no arbitrageur is allowed to close.

Why can no one capture it? Because the same mechanism that motivates the investment guarantees the giveaway. In the Aghion–Howitt model of **creative destruction** — the workhorse account of how market economies actually grow — innovation is replacement, not accumulation: a better product destroys the rents of the installed incumbent. The innovator's reward is a *temporary* monopoly profit; the free entry that makes the reward attainable also makes it perishable. Imitation, patent expiry and competitive entry then transfer the surplus to consumers as lower prices and better goods. Competition is the enforcement arm of the donation. The machine does not ask its winners to be generous; it forces the handover.

Here is the paradox the audit must hold honestly: the force that delivers 98 percent of the value to society is the same force that

starves the next round of its funding. This is not a contradiction — it is a dial. Patent length, antitrust intensity, open science and public research budgets are society's tuning of that dial, trading a little more private capture for a little more fuel. The evidence says the dial is set far below the optimum: returns of 30–70 percent do not persist in systems that are investing enough.

The public balance sheet is part of the answer, and the evidence supports it: the Dallas Fed's estimates find that government-funded R&D has delivered high long-run productivity returns of its own — the under-supply is fixable through public funding as well as private incentive, and the measured return says both channels are pushing on an open door.

Now scale the stakes. The world now spends roughly **\$2.9 trillion a year** on research and development. Apply the Jones–Summers range and the world's R&D effort generates on the order of **\$14.5–38.6 trillion of social value every year** — a flow comparable to adding the entire US economy, or two, annually, from an input costing under 3 percent of world output. Of that flow, the private actors who financed it capture a small fraction. Every year the engine runs under-fuelled, the forgone surplus is measured not in the billions of the missing budgets but in the multiples they would have returned.

### ENSI ESTIMATE – THE WORLD'S R&D DIVIDEND

Global R&D  $\approx$  **\$2.9T/yr** (2025)  $\times$  \$5–13.30 of social benefit per dollar = **\$14.5–38.6 trillion per year** of social value from the world's research effort — society's highest-return asset class, and its most under-supplied. ENSI estimate • CALCULATIONS §4 • Jones & Summers parameters

## The tension: ideas are getting harder to find

The case for the engine would be simple if research productivity were constant. It is not — it has been falling for as long as we can measure it. Bloom, Jones, Van Reenen and Webb, working across every domain where inputs and outputs can both be counted, find the same pattern: **sustaining a constant rate of progress requires an ever-larger army of researchers**, with total growth held steady only because research *effort* scaled massively to compensate (Stanford, 2020).

# 18<sub>x</sub>

Researchers now needed to sustain the Moore's-Law doubling, vs early 1970s

Bloom-Jones-Van Reenen-Webb • Stanford

# -5.1%/yr

Average annual decline in aggregate US research productivity

Bloom-Jones-Van Reenen-Webb • Stanford

# 41<sub>x</sub>

Cumulative fall in research productivity since the 1930s

Bloom-Jones-Van Reenen-Webb • Stanford

Read the two findings of this analysis together and the tension resolves into a single sentence: **the binding constraint on capitalism's dominant value mechanism is the volume of effective research effort** — an input getting more expensive precisely as its measured social return sits at 30–70 percent. Movement III returns to this machine and asks what happens when AI agents cut the cost of the one input the engine is short of.

### WHAT ANALYSIS 04 ESTABLISHES

- 1 **Innovation is the dominant term in the audit.** \$5–13.30 of social benefit per \$1 of R&D at floor and baseline, \$20+ at the ceiling — no other documented expenditure category returns in this range (Jones & Summers).
- 2 **The giveaway is structural, not incidental.** Spillovers dominate business-stealing at the firm level; producers keep ~2 percent at the economy level. The mechanism is competition, imitation and patent expiry — not benevolence (Bloom-Schankerman-Van Reenen; Nordhaus).
- 3 **Ideas are getting more expensive.** 18x the researchers for the same cadence; research productivity down ~41-fold since the 1930s. The constraint is effort volume, and it tightens annually.
- 4 **The policy arbitrage is standing.** Public and private R&D both earn measured 30–70 percent social returns; funding the gap is the closest thing macroeconomics has to a free lunch (Movement III).

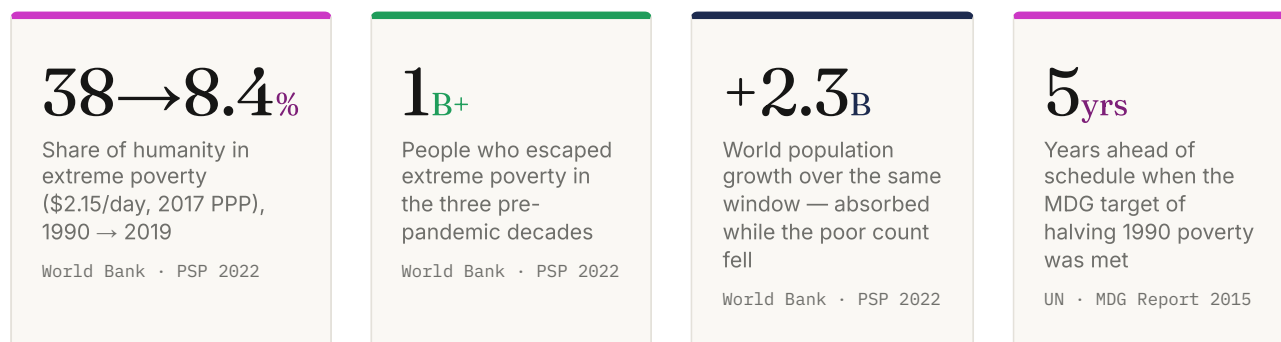
#### Read-across

Analysis 2 supplied the split (2.2% / 97.8%); this analysis showed why it under-fuels the machine that generates the surplus. Movement III prices what happens when agents enlarge the effective research base: a 20–30% cut in effective research cost, run through the Jones–Summers multiplier, is worth **+\$3–8T/yr** (ENSI estimate, CALCULATIONS §8).

## MOVEMENT I · THE AUDIT — ANALYSIS 05

## 05 The Poverty Event

In 1990, nearly **four humans in ten** lived in extreme poverty. By 2019, fewer than one in eleven did — while the planet added 2.3 billion people. It is the fastest mass improvement in human welfare ever measured, and this analysis records both its size and the honest allocation of the credit.



**B**egin with the baseline, because the baseline is what makes the event legible. In 1990, close to four in ten human beings lived below the extreme-poverty line — \$2.15 a day at 2017 prices, the threshold of bare subsistence. That figure was already an improvement on every earlier century; for most of recorded history, most of humanity lived at or near that line as a permanent condition. Poverty was not a policy failure. It was the default state of the species.


Then, in a single generation, the default broke. By 2019 the extreme-poverty rate had fallen to **8.4 percent**. More than a billion people crossed the line upward — and they did it while world population grew by roughly 2.3 billion. Hold that denominator in view, because it is what separates this event from every partial improvement before it: the global economy did not lift a fixed pool of the poor over a threshold. It absorbed **two billion additional human beings and still cut the absolute count of the destitute by over a billion**. Halving a rate while the denominator swells is a categorically harder feat than halving a count in a static population, and it is the feat the data record.

The accounting here is unusually clean, which is why this entry carries the audit's highest confidence grade. The World Bank has counted the same thing, the same way, at a consistent line, for a generation; the United Nations ran an independent flagship series over the same period on different survey infrastructure. The two agree on direction, magnitude and timing. Whatever else in this report rests on estimation, the poverty event does not.

What remains is the part an honest audit cannot skip: **who and what produced it**. Three-quarters of the decline happened in one country, and that country's own auditors credit markets *and* state capacity, not markets alone. Trade was a central channel. And a welfare-weighted reading — what a dollar is actually worth to someone living on two of them a day — makes this the most valuable entry in the entire audit per dollar moved. Those are the next three pages.

## Three decades, one line, one direction

FIG. 5.1 – EXTREME POVERTY, SHARE OF WORLD POPULATION, 1990–2020



**The default state of the species ended in one generation.** Extreme poverty (\$2.15/day, 2017 PPP) fell from roughly 38% of humanity in 1990 to 8.4% in 2019; COVID-19 produced the first sustained uptick in a generation (coral). World Bank, *Poverty and Shared Prosperity 2022*.

Two independent flagship series carry the same story. The Millennium Development Goal, set in 2000, aimed to halve the 1990 extreme-poverty rate by 2015; the UN's final accounting confirms the target was met **five years early**, with the rate more than halved and over a billion people out (UN, *MDG Report 2015*). The World Bank's series, built on different survey infrastructure, concurs on every order of magnitude. This matters for the audit because it retires the reflex that the number is an artefact of one institution's method: two instruments, one event.

### THE REVERSAL FOOTNOTE – RECORDED AT FULL WEIGHT


COVID-19 drove the first sustained rise in a generation: **9.3% in 2020**, more than 70 million people pushed back below the line, over 700 million in total. On current trends roughly **574 million people (~7%)** will remain in extreme poverty in 2030 — missing the 3% target. The escalator has slowed; the direction of three decades has not. World Bank · PSP 2022

“

*During the three decades that preceded its arrival, more than 1 billion people escaped extreme poverty.*

WORLD BANK · POVERTY AND SHARED PROSPERITY 2022 · VERIFIED AGAINST SOURCE

## The China honesty — and the trade channel



■ China's share of the entire global reduction in extreme poverty since 1978  
World Bank · Four Decades of Poverty Reduction in China

**Close to three-quarters of the entire global decline happened in one country.** The World Bank's dedicated study puts the count at close to 800 million fewer Chinese in extreme poverty since reform began in 1978 — roughly 75 percent of the world total, with 770 million lifted above the national line. It is the single largest anti-poverty event in human history, and it is not a story a market fundamentalist gets to tell straight.

The Bank's own attribution names **two pillars, not one**: rapid, broad-based growth unleashed by market reform — the de-collectivisation of agriculture, the opening to trade, the licensing of private enterprise — *and* a state with the capacity to build the roads, schools and clinics, and to target resources at the residual poor. An audit that credited the whole 800 million to laissez-faire would be as dishonest as one that credited it to central planning. The correct entry reads: the largest poverty exit on record was produced by **introducing market mechanisms into a society that also possessed unusual state capacity** — and the two together did what neither had done alone in the preceding centuries.

The channel that carried the event outward was trade. The joint WTO–World Bank study of trade and poverty is explicit that openness was central to the billion-person exit: connecting poor-country producers to world demand

### ~800M

fewer extreme poor in China since 1978 (\$2.15 line)

### 770M

lifted above China's national poverty line

converted rich-world purchasing power into poor-world wages at planetary scale, and developing economies' share of world exports rose across precisely the decades their poverty collapsed. The value-chain evidence sharpens the mechanism: joining cross-border production networks raises productivity, wages and formal employment in developing economies — it is how a Vietnamese or Bangladeshi worker comes to be paid something closer to the world price for their labour rather than the local one.

Independent measurement agrees. Sala-i-Martin's reconstruction of the entire world income distribution from 1970 to 2000 — built from national accounts and within-country distributions rather than poverty surveys — found the same double movement: a collapsing world poverty count alongside falling global inequality between individuals, driven by the growth of poor Asia. Different instrument, same event, same decades.

One more honesty, from the source itself: \$2.15 a day is a floor, not a comfort line. The World Bank notes that over 3 billion people still live below the higher lines used for middle-income countries. The event is real and enormous; the work is nowhere near finished — and the next page shows why, in welfare terms, every further dollar at the bottom counts more than anywhere else in the economy.

What a dollar is worth at two dollars a day

GDP arithmetic treats every dollar the same. Welfare arithmetic does not — and the difference is the reason the poverty event outranks its modest weight in world output. Under the standard log-utility assumption that anchors the Jones–Klenow welfare framework, the value of a marginal dollar falls in proportion to the income of the person receiving it. A dollar reaching someone at \$2 a day is therefore worth on the order of **10–30 times** the same dollar at rich-world incomes (Jones–Klenow method; ENSI application). The event this analysis audits moved income to exactly the people for whom income is most valuable.

Priced as a flow: lifting a billion people from roughly \$1.90 to \$5.50 a day equals about **\$1.3 trillion a year** of income now reaching the world's poorest (ENSI estimate). Against a \$110 trillion world economy that is barely one percent — and that is precisely the point. Weighted by what the dollars are worth to the people who received them, the smallest major entry in the GDP ledger becomes one of the largest entries in the welfare ledger. Any audit of what capitalism has paid society that stopped at unweighted dollars would miss it.

| LIVING STANDARD      | APPROX. DAILY INCOME | WELFARE WEIGHT OF A MARGINAL DOLLAR |
|----------------------|----------------------|-------------------------------------|
| Extreme poverty      | \$2/day              | ≈30x the reference dollar           |
| Post-exit poor       | \$6/day              | ≈10x                                |
| Global middle        | \$20/day             | ≈3x                                 |
| Rich-world reference | \$60/day             | 1x (reference)                      |

Log-utility weighting: the marginal value of a dollar scales with 1/consumption, shown relative to a \$60/day reference consumer. **ENSI illustration** of the Jones–Klenow method — the source of the 10–30x range used in the text.


WHAT ANALYSIS 05 ESTABLISHES

- 1 The event is real and twice-verified.** Two independent flagship series — World Bank and UN — concur on a fall from ~38% to 8.4% of humanity, 1990–2019 (very high confidence).
- 2 The denominator makes it historic.** The world absorbed 2.3 billion additional people while cutting the count of the destitute by more than a billion; the MDG halving target arrived five years early.
- 3 The credit is shared, honestly.** China accounts for ~75% of the decline, and its own auditors attribute the result to market reform *plus* state capacity — a finding to be quoted whole, never in half.
- 4 Trade was a central channel.** WTO–World Bank analysis ties the billion-person exit to openness; the value-chain evidence shows the wage mechanism at work.
- 5 Welfare-weighted, this is the audit's most valuable entry per dollar.** At 10–30x weighting, ~\$1.3T/yr flowing to the poorest (ENSI estimate) outweighs far larger flows higher up the distribution.

## MOVEMENT I • THE AUDIT — ANALYSIS 06

## 06 The Welfare Multiplier

GDP counts what people buy. It does not count how long they live to enjoy it. Price the mortality, the leisure and the equality alongside the consumption, and the welfare delivered by modern growth has been running roughly **1.5 times** the measured number — the audit's largest correction, and it points upward.



**I**ncome is a means. This audit would close hollow if the thirty-fold enrichment of Analysis 1 and the poverty exit of Analysis 5 had not converted into the things people actually want — years of life, healthy children, literate minds, hours not worked, the disappearance of routine catastrophe. It did convert, and this analysis prices the conversion. The headline finding: **measured growth has systematically understated the welfare delivered**, and by a wide margin.


The most direct statement of the multiplier comes from Jones and Klenow's welfare accounting, the most rigorous attempt yet to put consumption, leisure, inequality and mortality into a single consumption-equivalent number. Across the countries where the calculation can be run over time, **welfare growth averaged 3.1 percent a year against income growth of 2.1 percent** between the 1980s and the mid-2000s — welfare compounding half again as fast as GDP, chiefly because people were living longer and the extra years are worth a great deal (Jones & Klenow, AER 2016).

The physical fact underneath the accounting is the largest welfare gain any statistic captures: over the market era, life expectancy in the frontier economies roughly **doubled** — from the thirties and forties of 1900 to the high seventies and eighties today — and the poorest countries have been converging on the frontier faster in health than in income. None of that doubling is ever purchased in a shop, so almost none of it appears in GDP. A statistic built to track production will faithfully record the pharmaceutical's price and entirely miss the decades of life it added.

The pages that follow do the pricing in three steps. First, the master relationship — the **Preston curve** — and its most optimistic property: it keeps shifting up. Second, the mechanics: what actually drove mortality down, and how the Jones–Klenow method converts longevity, leisure and equality into consumption-equivalents, with France as the worked example. Third, the multiplication: what the welfare wedge does to the audit's headline counterfactual flow.

## The Preston curve — and why it keeps shifting up

FIG. 6.1 — THE PRESTON CURVE, SCHEMATIC: INCOME AND LIFE EXPECTANCY ACROSS COUNTRIES



**Income moves societies along the curve; knowledge moves the curve itself.** At any given income, people in later decades live substantially longer than people at the same income did earlier. Schematic after Preston (1975); the original curves are estimated cross-country relationships, redrawn here for exposition.

Samuel Preston's 1975 finding is the master relationship of this analysis: plot national income per head against life expectancy across countries, and the points trace a strong positive curve that **risers steeply among the poorest countries and flattens among the richest**. Money buys life most efficiently exactly where there is least of both — one more reason the poverty event of Analysis 5 is the audit's most welfare-dense entry.

But the deeper finding — the one that carries this analysis — is what happened to the curve over time. It did not sit still. **The entire curve shifted upward across the twentieth century**: at any given level of income, people lived substantially longer in later decades than earlier ones. Two forces, then, not one. Income growth moved societies *along* the curve; the diffusion of knowledge — germ theory, sanitation, antibiotics, vaccines, public health — moved the whole curve *up*. And both forces were downstream of the same market-driven enrichment: the growth paid for the science, built the water systems and manufactured the vaccines at prices poor countries could eventually afford. The doubling of life expectancy is the compound interest on both accounts at once.

**Preston 1975** WHO Bulletin reprint

**NBER w11963** Cutler-Deaton-Lleras-Muney

**OECD** How Was Life? 1820-present

**UNDP** Human Development Reports 1990 • 2023-24

## Pricing a longer life: the Jones–Klenow correction

### 67%

France's GDP per person vs the US, 2005

### 60%

France's consumption per person vs the US

### 92%

France's consumption-equivalent welfare vs the US

### ≈+10pp

each from lower mortality, lower inequality, higher leisure

What drove the mortality collapse is settled work. Cutler, Deaton and Lleras-Muney's decomposition finds no single cause but a compounding stack: **rising incomes** that bought food, housing and safety; the **germ theory of disease** and the public-health revolution it enabled — clean water, sanitation, vaccination; then the pharmaceutical era — antibiotics and the treatment of chronic disease. The gains landed first and largest among children, which is why the era's most transformed statistic is not average income but the probability of surviving to adulthood. Deaton's Nobel synthesis names the whole arc the **Great Escape**: the same process that freed humanity from material want freed it, with a lag, from early death — unevenly, he insists, and with stragglers the escape has not yet reached.

Pricing those gains is the contribution of Jones and Klenow. Their method asks a precise question: what fraction of one country's consumption, keeping its own leisure, mortality and inequality, would deliver the same expected utility as life in another? The answer is a **consumption-equivalent welfare** number that can be compared across countries and decades — the standard economics of expected utility, pointed at the biggest question there is.

The worked example rearranges the league table. France's GDP per person in 2005 was just 67 percent of the US level, its consumption only 60 percent. But the French lived around

**80 years against 77**, worked far fewer hours (535 per person per year against 877), and spread consumption more equally. Lower mortality, lower inequality and higher leisure each add roughly ten percentage points — lifting France from 60 percent of US living standards to **92 percent**. A third of the apparent gap was an artefact of the yardstick.

The correction is honest in both directions — which is what makes it credible. For many poor countries, the adjustment runs *downward*: shorter lives and higher inequality mean welfare below measured income, and welfare is more dispersed across countries than income is. The metric does not flatter growth; it prices it. And over time the price runs high: welfare growth of **3.1 percent a year against 2.1 percent** measured income growth — the ~1.5x multiplier this analysis is named for — because falling mortality is worth so much that it dominates the correction wherever health is improving.

“

*Rather than looking like 60 percent of the US value, as it does based solely on consumption, France ends up with consumption-equivalent welfare equal to 92 percent of that in the United States.*

JONES & KLENOW · BEYOND GDP? WELFARE ACROSS COUNTRIES AND TIME · AER 2016 · VERIFIED AGAINST SOURCE

## Two hundred years of welfare, priced

Now run the multiplier through the audit's headline. Analysis 1 established a counterfactual flow of roughly **\$150–180 trillion a year** of output above the pre-industrial baseline. That is a GDP-metric number — it counts the consumption and misses the lives. Apply the Jones–Klenow wedge, under which welfare-relevant growth has historically run about 50 percent above measured growth where mortality was falling, and the *welfare* value of modern growth lands at the equivalent of **\$225–270 trillion a year of consumption** — an order-of-magnitude estimate, labelled as such, but built from the most careful welfare accounting the literature has produced.

### ENSI ESTIMATE – THE WELFARE-ADJUSTED COUNTERFACTUAL

\$150–180T/yr counterfactual output flow (Analysis 1)  $\times$   $\approx 1.5$  welfare-to-consumption wedge (Jones–Klenow)  $\approx$  **\$225–270 trillion/yr consumption-equivalent**. The largest single line in the audit — and a floor, since it still excludes most of the free-goods surplus of Analysis 3. ENSI estimate · CALCULATIONS §6

The two-century record supports the direction of the correction everywhere we can measure it. The OECD's *How Was Life?* project, assembling global well-being data since 1820, finds the market era is also the era in which real wages, height — the biologist's proxy for childhood nutrition — literacy and schooling went from elite privileges to near-universal norms. And the UNDP's Human Development Reports bracket the modern arc: the first report, in 1990, was written precisely as a corrective to income-centric accounting — people as the real wealth of nations — and the 2023–24 edition's long-run series records a world higher on every component of the index, income, schooling and longevity alike, than the world the first edition measured, recent setbacks included.

### WHAT ANALYSIS 06 ESTABLISHES

- 1 **Life expectancy roughly doubled over the market era** — the single largest welfare gain any statistic captures, and one GDP records only incidentally (Preston; OECD).
- 2 **The Preston curve shifted up, not just along**. Income growth and knowledge diffusion compounded — and both run downstream of the same enrichment that funded the science and the sanitation.
- 3 **Welfare growth  $\approx 1.5$ x measured growth** — 3.1% vs 2.1% a year in the Jones–Klenow accounting — because added years of life are worth so much once priced.
- 4 **The correction is honest in both directions**. France rises from 60% to 92% of US living standards; high-mortality, high-inequality countries fall below their measured income. This is a yardstick, not an advocacy tool.
- 5 **Priced properly, modern growth is worth \$225–270T/yr consumption-equivalent** (ENSI estimate) — and the audit's headline numbers remain floors, not ceilings.

### Next: the debit column

Analyses 4–6 close the credit side of the audit: the engine, the exit, the multiplier. Analysis 7 turns the page to the debits — concentration, displacement and inequality — presented at full strength before any netting, because an audit that reports only credits is a brochure.

## MOVEMENT I · THE AUDIT — ANALYSIS 07

## 07 The Honest Debits

An audit that reports only the credits is not an audit but a brochure. Before this report nets anything, it books the four largest entries in the red column — at full strength, in their authors' own numbers, with nothing softened.

### 52%

Share of global income taken by the richest 10% of the world's population; the poorest half earns 8.5%

World Inequality Lab, 2022

### 2.4M

Upper estimate of US jobs lost to Chinese import competition, 1999–2011 — with decade-long local scarring

Autor, Dorn & Hanson

### 67%

Average US markup over marginal cost by 2014 — up from 18% in 1980, a three-and-a-half-fold rise

De Loecker & Eeckhout, NBER 2017

### 4%

The young-firm growth-skewness premium by 2007 — collapsed from 16% in 1999 and still falling through 2011

Decker et al., NBER 2015

The preceding six analyses assembled the credit side of the ledger: a \$150–180 trillion annual counterfactual flow, a ~98 percent surplus giveaway, a billion-person poverty exit, trillions of uncounted digital surplus, a doubled life expectancy. A reader is entitled to suspect an audit that finds only good news. So this analysis inverts the burden. The four debits below are the strongest documented entries against the machine, drawn from the same library and stated the way their authors state them — **not the way the machine's defenders would prefer them stated**.

They are: the concentration of income at the top of the global and national distributions; the concentrated, slow-healing job losses that trade inflicted on identifiable American communities; the four-decade drift of market power visible in markups and returns on capital; and the thinning of the high-growth entrepreneurial tail that does the machine's creative destruction. Each is real. Each is measured by economists with no brief for capitalism's defence — the World Inequality Lab, the labour economists who named the China Shock, the industrial-organisation team that built the markup series.

Two of the four, we will argue, are questions about **who holds the surplus**; two are **maintenance faults that degrade tomorrow's surplus**. But that classification comes later. Pages 39 and 40 present the red ink uninterpreted. Page 41 does the netting — and shows the arithmetic that keeps the audit, decisively, in the black.

#### HOW TO READ THIS SPREAD

Coral marks the debit column throughout Analysis 07 — the only pages of this movement where the alert colour carries the argument. Nothing on pages 38–40 is netted, discounted or answered. The response begins on page 41 and continues through Analyses 08 and 09.

## ANALYSIS 07 · THE DEBIT COLUMN, ENTRIES ONE AND TWO

**Debit 1 — Concentration at the top**

**T**he distributional headline of the market era is not subtle. The World Inequality Report's flagship finding, in its own words: "The richest 10% of the global population currently takes 52% of global income, whereas the poorest half of the population earns 8.5% of it" (World Inequality Lab, 2022). Translated to persons: an average adult in the global top decile earns about **\$122,100 a year**; an average adult in the poorer half of humanity makes about **\$3,920** — a thirty-one-fold gap between the top tenth and the bottom half, inside the most prosperous economy the species has ever run.

The national series sharpen the charge. Piketty and Saez's reconstruction of US top income shares from tax records shows a U-shape: the top decile's share of national income was compressed hard through the mid-century

decades, then **rebounded after 1980** back toward its pre-war heights (Piketty & Saez, Berkeley). Whatever mid-century forces flattened the distribution — war, tax policy, strong unions, financial regulation — the post-1980 policy mix let concentration reassert itself. The market era did not merely tolerate a widening gap in the rich world; on this evidence it produced one, over exactly the decades its defenders celebrate.

FIG. 07-A — THE GLOBAL INCOME SPLIT, 2021



**Half of humanity shares one-twelfth of the income.** Shares of global income at purchasing power parity; middle-40% share is the arithmetic remainder. Source: World Inequality Lab, *World Inequality Report 2022*.

**Debit 2 — The China Shock: concentrated, slow-healing losses**

Trade's aggregate dividend was booked in Analysis 3's credit column at \$2.1 trillion a year for the United States alone. Here is its gross cost, from the economists who counted it. The surge of Chinese imports between 1999 and 2011 eliminated an estimated **985,000 US manufacturing jobs directly**; adding input-output linkages raises the toll to 2.0 million economy-wide, and counting local labour-market spillovers pushes the range to **2.0–2.4 million** (Autor, Dorn & Hanson).

The number is not the sting. The sting is what happened next — which was, for a decade, approximately nothing. The textbook promised displaced workers would reallocate to expanding sectors. They largely did not. The authors' verdict deserves quoting at length, because it overturned the field's received wisdom:

“*Adjustment in local labor markets is remarkably slow, with wages and labor-force participation rates remaining depressed and unemployment rates remaining elevated for at least a full decade after the China trade shock commences.*”

**AUTOR, DORN & HANSON** · "THE CHINA SHOCK," ANNUAL REVIEW OF ECONOMICS

Exposed workers suffered greater job churning and **reduced lifetime income**; the pain pooled in the specific counties of the Midwest and Southeast where the exposed industries clustered. A national surplus coexisted with regional devastation — and the compensation the theory assumed would follow largely never arrived.

## ANALYSIS 07 • THE DEBIT COLUMN, ENTRIES THREE AND FOUR

**Debit 3 — The market-power drift**

**T**he whole case for the machine rests on one property: competition forces the surplus through to consumers. Debit three is the four-decade evidence that this forcing mechanism is wearing. De Loecker and Eeckhout, measuring markups across US publicly traded firms, document a rise in the average markup from **18 percent above marginal cost in 1980 to 67 percent in 2014** — a three-and-a-half-fold increase, concentrated in the upper tail of the firm distribution rather than spread evenly (De Loecker & Eeckhout, NBER 2017). Every point of markup is a point of surplus that stops escaping to the public. On the audit's own logic, this series is an attack on its central asset: the 98 percent giveaway is


**Debit 4 — The thinning of the experiment**

The machine creates value through entry — young firms running the up-or-out experiments that occasionally produce the next transformative company. Debit four says the experiment is losing its tail. Decker and co-authors document that US business dynamism has declined for decades, and that **the character of the decline changed around 2000**. In 1999, the firm at the 90th percentile of employment growth grew about 31 percent faster than the median, and that 90-50 differential ran 16 percent larger than the 50-10 spread below it. **By 2007 the skewness premium had collapsed to 4 percent** — and it kept falling through 2011 — driven by a falling share of young firms and a falling propensity of the young firms that do enter to be high-growth (Decker, Haltiwanger, Jarmin & Miranda, NBER 2015). Fewer experiments, and tamer ones: the exact opposite of what a healthy Schumpeterian engine looks like from the outside.

not a law of nature, and here is the gauge showing it can erode.

The White House Council of Economic Advisers' indicators point the same way. Revenue shares of the top 50 firms rose in most broad sectors between 1997 and 2012 — transportation and warehousing by 11.4 percentage points, retail by 11.2. And the returns tell the sharpest version: the **90th-percentile firm's return on invested capital now runs at more than five times the median**, a ratio the CEA notes was closer to two just a quarter-century earlier (CEA, 2016). Super-normal returns that persist are the signature of positions that rivals cannot contest.

FIG. 07-B — THREE GAUGES, FOUR DECADES OF DRIFT



**Coral marks the deteriorated reading of each gauge.**

Markups and top-firm returns rose; the young-firm growth premium collapsed. Sources: De Loecker & Eeckhout (NBER 2017); CEA (2016); Decker et al. (NBER 2015).

**ONE FAULT, SEEN FROM TWO SIDES**

The markup rise and the dynamism decline are plausibly the **same** degradation observed in different data: when incumbents' positions harden, challengers enter less and grow less. Debits 3 and 4 share a cause — and, as Analysis 08 will show, a toolkit.


## ANALYSIS 07 • THE NETTING

**Real, bounded, concentrated — and one to two orders of magnitude smaller than the credits**

Now the netting, with the debits conceded in full. Three things are true of all four entries. They are **real** — measured, replicated, undisputed in direction. They are **bounded** — none reverses or approaches the scale of the credit entries. And they are **concentrated** — which is precisely why they dominate the politics: a \$2.1-trillion national dividend is diffuse and invisible, while a closed factory has an address.

Classify them and the ledger clarifies. Debit 1 is a question about **who holds the surplus**, not about whether it exists: the 52 percent the global top decile takes is a share of an income pool the machine multiplied thirty-fold — a distribution problem inside a creation success, treated in full in Analysis 09. Debit 2 is a gross loss already netted inside a credit entry that remains roughly **ten-to-one positive**: against the \$2.1 trillion annual US trade dividend, the China Shock's adjustment costs run on the order of \$0.2 trillion a year (ENSI estimate, from the ~10:1 net in this report's calculations) — real, unequally borne, and an indictment of the missing compensation layer rather than of exchange itself. Debits 3 and 4 are not destroyed value at all but **maintenance faults**: warning-gauge readings that tomorrow's giveaway will shrink if the competitive machinery is left unserviced. Nowhere in the 106-document library is there evidence of a failure in the machine's core function — that innovation stopped returning multiples, prices stopped aggregating information, or entry stopped creating jobs. The faults sit in the enforcement and access layers: the wear surfaces.

FIG. 07-C – DEBIT MAGNITUDES AGAINST CREDIT MAGNITUDES, \$T PER YEAR



Bar lengths are log-scaled; on a linear scale the coral bar would be invisible — which is itself the finding. Credit flows are annual magnitudes from Analyses 1–4 (BOE aggregations labelled in those sections; ENSI estimates). The China Shock bar is an ENSI estimate consistent with the ~10:1 trade net. Debits 1, 3 and 4 do not appear as bars because they are distributional shares and trend gauges, not annual value destruction — the honest reason, not a convenient one.

The netting statement, then: **the debits are real and the credits run one to two orders of magnitude larger** — and the two debit entries that most threaten the future (markups, dynamism) are exactly the ones with known, tested fixes. Which raises the question the next analysis answers: if the faults are fixable, what is the repair manual?

WIL World Inequality Report 2022

Berkeley Piketty &amp; Saez, US top shares 1913–1998

ADH The China Shock, Annu. Rev. Econ.

NBER De Loecker &amp; Eeckhout w23687

CEA Benefits of Competition, 2016

NBER Decker et al. w21776

## MOVEMENT I · THE AUDIT — ANALYSIS 08

## 08 The Maintenance Thesis

Every debit on the preceding spread is an engineering fault with a documented fix. Read properly, the market-failure literature is not the machine's obituary. It is the machine's specification sheet — and its repair manual has been on the shelf for decades.

Start with the strongest theoretical objection, because the thesis stands or falls on absorbing it. Joseph Stiglitz's work on the invisible hand demonstrates that under imperfect and asymmetric information, with incomplete markets, competitive outcomes are generically **not constrained-efficient** — the assumptions of the First Welfare Theorem fail pervasively, so market results can in principle be improved upon (Stiglitz, NBER w3641). This is decisive against *laissez-faire as a theorem*. It is not decisive against the machine, for a reason engineers will recognise: a specification that states the operating conditions under which a system performs — and the conditions under which it degrades — is the most useful document a system can have. Stiglitz tells us **where the faults will occur**: wherever information is asymmetric and costs fall on parties outside the transaction. That is a map of the maintenance schedule, not a demolition order — especially since the alternative coordination technology, a planner with better information than the price system, is precisely what the calculation debate showed cannot exist.

Next, the repair manual. Elinor Ostrom's Nobel lecture documents a large empirical class of commons and externality problems — the canonical market failures — solved in practice by **polycentric, self-governing institutions**: neither market nor state, but nested layers of local monitoring, graduated sanctions and cheap conflict resolution (Ostrom, Nobel Lecture, 2009). The correction layer for price failures, her fieldwork shows, is itself usually decentralised. And at the level of the firm, Hart and Zingales argue that the objective function is adjustable from inside: companies should maximise **shareholder welfare, not market value** — because shareholders are people with ethical and social preferences, the corporate machine can internalise harms its owners actually care about without any change to its ownership structure (Hart & Zingales).

The honest limit of the thesis is political, and the library states it without flinching:

“

*Market concentration can easily lead to a "Medici vicious circle," where money is used to get political power and political power is used to make money.*

LUIGI ZINGALES · TOWARDS A POLITICAL THEORY OF THE FIRM, NBER W23593

### THE CLAIM, PRECISELY

Market failures are real, local, and correctable through institutions. The rational response to a worn brake pad is a brake job — not a referendum on braking.

## ANALYSIS 08 · THE DASHBOARD

**The gauges: five instruments on the machine's dashboard**

If capitalism is infrastructure, it should be monitored like infrastructure — continuously, on named gauges, with thresholds that trigger service. The library supplies exactly five. Each has a published instrument, a current reading, and a toolkit that has worked before. Two are in the red; two are amber; one is the meta-gauge that decides whether the other four ever get serviced.

| GAUGE                        | INSTRUMENT                                                              | CURRENT READING                                                                                           | STATUS                    | SERVICE TOOLKIT                                                                               |
|------------------------------|-------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------|---------------------------|-----------------------------------------------------------------------------------------------|
| <b>Markup trend</b>          | De Loecker–Eeckhout markup series across US listed firms                | 67% over marginal cost (2014), from 18% (1980); rise concentrated in the upper tail                       | RED · rising              | Antitrust enforcement, merger control, entry-barrier removal                                  |
| <b>Entry rate</b>            | Census firm-entry series (CEA synthesis)                                | Secular decline in the US firm-entry rate over recent decades; concentration rising in most broad sectors | RED · falling             | Occupational-licensing reform, entry-cost demolition, finance access                          |
| <b>Young-firm skewness</b>   | Decker et al. 90-50 vs 50-10 growth-differential premium                | 16% (1999) → 4% (2007), still falling through 2011; fewer young firms, tamer ones                         | RED · collapsed           | Shares gauges 1–2 toolkit, plus venture-finance depth                                         |
| <b>Allocative dispersion</b> | Hsieh–Klenow within-industry productivity spread                        | 90/10 revenue-productivity ratio ~6–7× in China and India vs under 2× in the US                           | AMBER · reform-responsive | Factor-market liberalisation, financial deepening                                             |
| <b>Governance</b>            | Worldwide Governance Indicators; capture risk per Zingales' Medici test | Capture risk co-moves with concentration — the meta-gauge on gauge 1                                      | AMBER · watch             | Lobbying transparency, polycentric oversight, institutional stress-tests each political cycle |

Sources: De Loecker & Eeckhout (NBER 2017); CEA (2016); Decker et al. (NBER 2015); Hsieh & Klenow (QJE 2009); World Bank WGI; Zingales (NBER 2017).

Read the dashboard as a whole and one pattern jumps out: **the red gauges are all in the enforcement and access layers** — competition, entry, dynamism — while the machine's core functions (innovation returns, price signalling, job creation by entrants) show no documented degradation anywhere in the library. That is the profile of a maintenance backlog, not a design failure. It is also the profile of a system whose faults are coupled: the markup drift feeds the entry decline, the entry decline thins the skewness tail, and unchecked concentration eventually reaches the governance gauge — the one fault that, if it trips, disables the repair crew itself. Hence the maintenance ordering: **service gauge 1 first, and instrument gauge 5 permanently**.

## ANALYSIS 08 • OBJECTIONS, AT FULL STRENGTH

**The steelman, answered line by line**

The maintenance thesis earns its name only if it survives the strongest versions of the objections — stated the way their proponents state them, then answered or conceded. Six rows; two concessions inside the answers; no objection paraphrased down.

| THE OBJECTION, AT FULL STRENGTH                                                                                                                                                | THE MAINTENANCE RESPONSE                                                                                                                                                                                                                                                                                            |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Under asymmetric information and incomplete markets, market outcomes are <b>generically inefficient</b> — the welfare theorems fail pervasively, not at the margin (Stiglitz). | <b>Concede the theorem.</b> It licenses targeted instruments — disclosure rules, insurance design, externality pricing — layered on the price system. It does not supply a substitute signalling layer; no planner has less of an information problem than the market it would replace (Hayek).                     |
| Markets mishandle <b>commons and externalities</b> ; unpriced harms accumulate — carbon being the largest (Ostrom; Stiglitz).                                                  | <b>Concede, and cite the fieldwork:</b> Ostrom documents commons governed successfully by polycentric self-organisation. Carbon is the largest open ticket in the maintenance queue — an unpriced liability, not a refutation of pricing.                                                                           |
| <b>Market power is rising:</b> markups 18%→67% over cost since 1980; top-firm returns pulling five-fold away from the median (De Loecker–Eeckhout; CEA).                       | The series is real and is this report's most serious warning light. The toolkit is the oldest in the book — enforcement, merger control, entry-barrier demolition — and it has worked before: airline deregulation alone returned ~\$28B a year to consumers.                                                       |
| <b>Dynamism is dying:</b> the high-growth young-firm tail has thinned since 2000; the experiment engine is slowing (Decker et al.).                                            | Substantially downstream of the market-power fault, and it shares the toolkit — plus licensing reform and finance access. A coupled fault serviced by the same repair.                                                                                                                                              |
| <b>Within-country inequality rose</b> across the rich world after 1980, under exactly the policy mix the machine's defenders favoured (Chancel & Piketty; Piketty & Saez).     | True, and not netted away here: Analysis 09 treats it in full. It is a question about the distribution of a created surplus — with fiscal and predistributive instruments — not evidence of value destruction.                                                                                                      |
| <b>The winners will block the fixes:</b> concentration buys political power, which protects concentration — the Medici vicious circle (Zingales).                              | <b>The strongest objection, and the least answerable by economics alone.</b> The response is institutional: keep the governance gauge permanently instrumented, keep enforcement polycentric so no single captured node disables it, and widen the corporate objective itself (Hart–Zingales' shareholder welfare). |

What the maintenance framing buys is not comfort but **jurisdiction**. "Is capitalism good?" is a referendum — unanswerable, unfalsifiable, permanent. "Which gauge is in the red and what is the service procedure?" is a work order. Every objection above converts into a line on that work order, which is the entire point: the critics' evidence, taken at full strength, specifies the repairs rather than the demolition. One question survives the conversion, though — the one no gauge settles: even a well-maintained machine distributes its output somehow. Who, over two centuries, actually got the Great Enrichment?


**NEXT**

Analysis 09 answers with the longest-run dataset in the library: two hundred years of world income distribution, and two curves moving in opposite directions.

## MOVEMENT I · THE AUDIT — ANALYSIS 09

## 09 The Distribution Question

Who got the Great Enrichment? Two hundred years of data return a double answer: the gap **between** countries rose for 160 years and then began to close, while the gap **within** countries compressed for 70 years and then reopened. Most political argument quotes one curve and hides the other.



The authoritative source is Chancel and Piketty's reconstruction of the world income distribution from 1820 to 2020, built on the World Inequality Database. Its headline is sobering: global inequality "has always been very large," with the top decile's share oscillating in a 50–60 percent corridor across the entire period. But the level hides the mechanics, and the mechanics are the finding. Global inequality has **two components** — the gaps between countries and the gaps within them — and for the last century they have moved in opposite directions:

“

*Within-countries inequality dropped in 1910–1980 (while between-countries inequality kept increasing) but rose in 1980–2020 (while between-countries inequality started to decline).*


CHANCEL & PIKETTY · GLOBAL INCOME INEQUALITY 1820–2020, JEEA 2021

Hold both clauses at once, because each political camp owns exactly one. The globalising decades after 1980 produced the poverty collapse and the first sustained convergence between nations in modern history — **and** a widening of the gap inside most rich nations. The critic who cites only the second and the champion who cites only the first are each telling half the story. This analysis tells both halves, then asks what follows.

## ANALYSIS 09 · THE DECOMPOSITION

## Two curves, one question

FIG. 09-A – GLOBAL INEQUALITY DECOMPOSED: BETWEEN VS WITHIN COUNTRIES, 1820–2020 (SCHEMATIC)



The two components of global inequality have moved in opposite directions since 1910 — and swapped roles in 1980. Schematic rendering of the decomposition in Chancel & Piketty, *Global Income Inequality 1820–2020* (JEEA 2021); curve shapes indicate direction and turning points, not measured values.

The schematic makes the audit's distributional finding legible in one look. For 160 years, the dominant engine of global inequality was **geography** — which side of the industrialisation frontier your country sat on, an accident of birth no individual effort could overcome. The mid-century compression inside rich countries (1910–1980) coincided with the era's wars, taxation and welfare-state construction. Then 1980: the curves cross trajectories. As markets integrated, the poor giants grew faster than the rich world for the first time in the modern record — pulling the navy curve down — while the rules-mix inside rich countries let their internal gaps reopen, pushing the gold curve up. The result, Chancel and Piketty note, is that early 21st-century capitalism carries a level of global inequality comparable to that of 1910 — **but produced by a completely different mechanism, under a different set of rules**. And rules, unlike geography, can be rewritten.

ANALYSIS 09 • THE INCIDENCE

Who got the Great Enrichment

Decompose the last four decades by who gained, and the shape that emerges is the one made famous by elephant-curve readings of the same data: **large gains at the global middle, large gains at the global top, and a squeezed stretch in between.** The global middle is emerging Asia — the hundreds of millions who moved from farms into the world trading system and roughly doubled and redoubled their incomes. The global top is the rich world's capital owners and top earners, whose national income shares rebounded after 1980 (Piketty & Saez). The squeezed stretch is the rich world's lower-middle — the China Shock counties of Analysis 07, whose relative position stagnated while both their national elites and their overseas competitors advanced.

The convergence side of that story is the largest welfare event in the dataset, and it deserves its receipts. Extreme poverty fell from roughly **38 percent of humanity in 1990 to 8.4 percent in 2019**, with more than a billion people exiting while world population grew by 2.3 billion (World Bank, 2022). China alone accounts for close to **800 million** of the exits — about three-quarters of the global decline — produced, the World Bank is explicit, by market reform *plus* state capacity (World Bank, *Four Decades*). Vietnam's income per head rose from **\$231 to \$2,171** between 1985 and 2016 after the Doi Moi reforms (UNU-WIDER). Poland converged from the European

periphery to **62 percent of Western European income** by 2013, the continent's fastest climb (World Bank). These are not abstractions offsetting rich-world pain in a utilitarian ledger; they are the between-country curve of Fig. 09–A bending down, person by person.

| RECEIPT                                      | MAGNITUDE                               | SOURCE                   |
|----------------------------------------------|-----------------------------------------|--------------------------|
| Extreme-poverty rate, 1990→2019              | 38% → <b>8.4%</b>                       | World Bank PSP 2022      |
| China: exits from extreme poverty since 1978 | <b>~800M</b> (~75% of global)           | World Bank, Four Decades |
| Vietnam: GDP per capita, 1985→2016           | \$231 → <b>\$2,171</b>                  | UNU-WIDER 2018           |
| Poland: share of W-European income by 2013   | <b>62%</b> , Europe's fastest 1989–2013 | World Bank / Piatkowski  |

Now weigh the two curves against each other honestly. The within-country reopening is real, measured in shares of rich-world income. The between-country closing is measured in **escapes from destitution** — and on any welfare accounting that values a dollar more in the hands of someone at \$2 a day than at \$200 (the log-utility weighting of Analysis 5, worth 10–30×), the convergence gains dominate by an enormous margin. That is not a reason to dismiss the within-country curve; it is the reason the response to it must not sacrifice the between-country one. Which brings the audit to the policy fork that Analysis 07 left open.

## ANALYSIS 09 · THE POLICY FORK

**The compensation layer: the fix is transfers and mobility, not autarky**

The China Shock has two candidate remedies, and they are not symmetric. The first — close the border, repatriate the supply chains — cancels the roughly **\$2.1 trillion annual dividend to protect against a ≈\$0.2 trillion adjustment cost** (ENSI estimate, Analysis 07): a ten-for-one destruction of national income that would also, applied globally, throw the between-country convergence machine into reverse against the very populations it just lifted. The second remedy accepts what Autor, Dorn and Hanson actually demonstrated. Their finding was never that trade fails to pay; it was that **the adjustment the textbook assumed does not happen on its own** — wages and participation stay depressed for a decade in exposed places. That is a finding about a missing policy layer, not about exchange. The fix it implies is a compensation layer funded out of the dividend it protects.

Even the authors of the inequality series point down this road. Chancel and Piketty's own forward-looking discussion concerns **alternative rules** — they model how fiscal revenue-sharing arrangements could significantly reduce global inequality — a redesign of the distribution machinery, not an abandonment of the production machinery. The two-century lesson of their data is that the top-decile share is set by rule-sets: it compressed under one (1910–1980) and reopened under another (1980–2020), around a production engine that ran throughout. Distribution is a policy output. Treat it as one.

- 1 **Keep the aggregate gain.** Openness stays; the dividend is the fund from which every repair is paid. A remedy that shrinks the surplus to rearrange it fails both curves at once.
- 2 **Pre-commit the compensation.** Adjustment assistance decided *after* the losses concentrate arrives a decade late — the decade ADH measured. Trigger it automatically on exposure, not on politics.
- 3 **Fund people and places.** Decade-long local scarring means mobility support alone is not enough; the evidence indicts both the worker's stranded skills and the county's stranded tax base.
- 4 **Instrument the within-country gauge.** Top-share telemetry (Piketty–Saez, WID) belongs on the maintenance dashboard of Analysis 08, watched continuously like the markup series.
- 5 **Fix distribution with distribution tools.** Taxes, transfers, predistribution and revenue-sharing act on the gap directly; tariffs act on it by burning the surplus. Match the instrument to the fault.

The audit's distributional verdict, then: the Great Enrichment was real and increasingly widely shared **between** nations; its sharing **within** nations is a live, worsening, and entirely policy-tractable fault. What the verdict still lacks is a control group — what the alternatives to this machine actually paid the societies that ran them. That is Analysis 10.

PSE/JEEA Chancel &amp; Piketty 1820–2020

WIL World Inequality Report 2022

ADH The China Shock

PIIE Hufbauer &amp; Lu, PB 17–16

World Bank PSP 2022 · Four Decades

UNU-WIDER Vietnam 2018

## MOVEMENT I • THE AUDIT — ANALYSIS 10

## 10 • The Systems Counterfactual

Analyses 1–9 priced what the market system delivered. This analysis prices what the alternatives delivered when actually tried — and the gap is the cleanest single piece of evidence in the file.

The standard critique compares capitalism with perfection and finds it wanting. An audit cannot work that way; the only defensible comparison is with the systems that actually competed with it. The twentieth century, at appalling human cost, twice split a single people on a single geography into two operating systems — Korea in 1948, Germany in 1945–49 — and within two generations the income gaps ran at **ten-to-one to twenty-to-one**. No other variable in social science moves an outcome by an order of magnitude while holding everything else constant.

The evidence arrives in four registers. The **theoretical root**: Hayek's demonstration of why a planning apparatus cannot, even in principle, do what a price system does. The **micro ledger**: the state-versus-private ownership literature. The **marginal measurement**: Hsieh and Klenow's misallocation accounting — socialism in miniature, priced per dose. And the **before-and-after panels**: Poland, Vietnam, Estonia and the Korean split.

These are histories, not randomised trials — but read together they control for continent, culture, colonial past, resources and starting income, and the gradient always runs the same way. The Fraser index states the cross-section bluntly: the freest quartile of economies produces **7.6 times** the GDP per head of the least-free quartile.

### 10–20:1

Same-people, same-geography income gaps opened by the Korean and German system splits

ENSI estimate • Maddison / natural experiments • CALC §1

### +30–50%

China manufacturing TFP gain from reallocating inputs to US-level allocative efficiency

Hsieh & Klenow • QJE 2009

### 62%

Poland's GDP per capita as a share of Western Europe's by 2013 — Europe's fastest 1989–2013

Piatkowski • World Bank 2015

### 9.4<sub>x</sub>

Vietnam GDP per capita, \$231 → \$2,171 (1985–2016), dated to the 1986 Đổi Mới reforms

UNU-WIDER 2018

## The root and the receipts: from the calculation debate to the ownership ledger

Why did every planned economy underperform, everywhere, under every leadership? The answer was written before most of the experiments began. Hayek's 1945 argument is the deepest single insight in this library: the economic problem is not allocation given full information — it is that the relevant knowledge exists only as **dispersed, contradictory, local fragments** in millions of separate minds. Knowledge of the particular circumstances of time and place — which machine has an hour of slack, which cargo went unsold — is fleeting and tacit; no statistical bureau can collect it before it is obsolete. Prices solve the problem as a telecommunications system: tin becomes scarcer — and it does not matter to any user why — the price rises, every user economises, every producer of substitutes expands, the whole world adjusts without central instruction. His companion essay completes the case: competition is a **discovery procedure** — the facts a market coordinates do not exist before the market process generates them, so no planner can substitute for it, even in principle.

This is why the socialist calculation debate was not settled by ideology but by bandwidth. The planner's failure is an information failure before it is an incentive failure: the machine that was supposed to replace prices had no input channel for the knowledge prices carry. The twentieth century then confirmed the theorem empirically, at continental scale, with the results priced in Analysis 1.

The micro ledger says the same thing at the level of the firm. Shleifer's survey of state

versus private ownership — the nearest thing economics has to an A/B test of administratively guided against price-guided production — finds **private, price-disciplined operation systematically outperforming state operation** across industries and countries: the state manager faces neither the signal (prices and profit) nor the consequence (loss and exit) that force costs down and quality up. The result survives sector by sector, from airlines to banks.

Djankov, Glaeser, La Porta, López-de-Silanes and Shleifer then generalised the question into a field — the new comparative economics. Instead of asking "market or state?" as theology, it measures institutional alternatives on outcomes: every society chooses a point on a trade-off between the costs of private **disorder** and the costs of public **dictatorship**, and the data reward arrangements that discipline both — private operation under rules, with the state as referee rather than owner. The question this audit inherits is never whether rules are needed, but which control structure processes the most information at the lowest cost of abuse.

“

*The twentieth century ran the experiment at continental scale — same people, same geography, different operating system — and the gap came out ten-to-one.*

ENSI · THE VALUE ENGINE — THE AUDIT'S CONTROL GROUP

AER 1945 Hayek · The Use of Knowledge in Society [01]

Hayek Competition as a Discovery Procedure [01]

NBER 1998 Shleifer · State versus Private Ownership [01]

NBER 2003 Djankov, Glaeser, La Porta, López-de-Silanes & Shleifer [01]

Socialism in miniature: the misallocation meter

Every market economy administers some of its allocation — subsidised credit, licensing, connected firms, state banks — and Hsieh and Klenow built the meter that reads the resulting damage. Their method compares the marginal products of capital and labour across manufacturing plants **within narrow industries**. In a frictionless market these converge: inputs move until the gaps close. Where administration overrides prices, the gaps persist — and their dispersion measures how far allocation has drifted from what the price system would have chosen.

In China and India, the 90th-to-10th percentile ratio of plant-level revenue productivity runs at a factor of **6–7, against under 2 in the United States**. Hypothetically reallocating inputs to US-level efficiency — no new technology or capital, just prices allowed to do their job — would raise manufacturing TFP by **30–50 per cent in China and 40–60 per cent in India**, with output gains approaching a factor of four once capital accumulation responds. Misallocation is one of the largest sources of forgone income on the planet.

FIG. 10.1 – THE PRICE OF OVERRIDING PRICES



Letting prices allocate is worth a third to a half of manufacturing productivity. Hypothetical TFP gains from reallocation to US-level allocative efficiency (bars to range midpoints). Hsieh & Klenow, QJE 2009 [07].

6–7x

90th/10th percentile plant revenue-productivity ratio, China & India — the dispersion prices would arbitrage away

<2x

The same ratio in the United States — the price-disciplined benchmark

~2%/yr

China's own TFP gains from reallocation, 1998–2004 — the reform decades running the repair in real time

Read the right way round, the meter is the systems argument in marginal form: every administrative override of the price signal is a small dose of the same disease that killed the planned economies; the full command economy is simply the 100 per cent dose. The dose-response curve is continuous — China's partial reforms bought partial recovery exactly over the years its markets were liberalising. The counterfactual is not a Cold-War memory; it is a live gradient every economy sits on today.

WHY THIS BELONGS IN THE AUDIT

Hsieh-Klenow turns "capitalism versus the alternatives" from a historical claim into a measurable, continuous variable: the further allocation drifts from prices, the larger the measured loss — a factor of two in manufacturing productivity at the doses China and India were running.

## Before and after: four societies that switched the machine on

The strongest evidence macro-history allows is the recorded before-and-after. Four panels, four decades, four continents' worth of excuses controlled away.

### Poland · 1989 →

periphery → 62% of Western Europe

From the collapse of communism to 2013, Poland grew more than any other economy in Europe and was the only EU economy to avoid recession in 2008–09. By 2013 its GDP per capita reached about **62 per cent of the Western-European level** — the fastest convergence on Western living standards in the continent's modern history.

**World Bank** Piatkowski · Poland's New Golden Age [15]

### Vietnam · Đổi Mới 1986 →

\$231 → \$2,171 GDP/capita, 1985–2016

A byword for war and famine chose to let markets and prices work. Income per head rose almost **tenfold in three decades**, alongside one of the fastest poverty declines ever recorded — achieved with comparatively low inequality.

**UNU-WIDER 2018** The Dragon That Rose From the Ashes [15]

### Estonia · 1992 →

Soviet collapse → digital frontier in ~15 years

The most radical post-Soviet transition: flat income tax, near-total trade liberalisation, hard money, mass privatisation — and, presciently, an early digital state. **Shock liberalisation, competently executed, compressed a generation of catch-up into fifteen years.**

**Heritage** Laar · The Estonian Economic Miracle [15]

### Korea · split 1948 →

one people, two systems, order-of-magnitude gap

The South — at roughly African income levels in 1960 — escaped the middle-income trap through **private-sector-led competition, export discipline and a climb into frontier R&D**, joining the high-income OECD within two generations. The North ran the control condition on the same peninsula.

**World Bank-KDI** Innovative Korea [15]

**ENSI** gap: CALC §1 [15][09]

Read together, the panels control for continent, culture, colonial history, natural resources and starting income — communist and post-colonial, Asian and European, resource-poor and war-ruined — and the one variable that tracks the outcome every time is **the adoption of market mechanisms under decent institutions**. That is as close to a controlled experiment as macro-history allows, and it points one way. The alternatives were not cheaper, fairer versions of the machine; they were the same populations earning a fraction of the income, and in every recorded case the switch — not the endowment — moved the number.

### WHAT THE ALTERNATIVES PAID

The systems counterfactual is the audit's floor: whatever debits capitalism carries (Analysis 7), the competing operating systems delivered order-of-magnitude less income to the same people on the same land — and the cost of partial substitution is still measurable, per dose, wherever administration overrides prices.

## MOVEMENT I · THE AUDIT – ANALYSIS 11

## 11 · The Agentic Bet

Goldman Sachs says AI adds ~7 per cent to global GDP. Acemoglu's bound says at most 0.71 per cent of TFP in ten years. Both use the same task data. The ten-fold gap is not a disagreement about coefficients — it is a bet on whether AI stays a **tool** or becomes an **actor**.

**T**he loudest macroeconomic argument of the decade looks like a duel of forecasts. Goldman Sachs estimates that generative AI could raise global GDP by roughly **7 per cent — about \$7 trillion** — and lift US productivity growth by almost **1.5 percentage points a year** over the decade after widespread adoption, with the equivalent of 300 million full-time jobs exposed globally. Acemoglu, applying Hulten's theorem to the same exposure data, bounds total-factor-productivity gains at **no more than 0.71 per cent over ten years** — and nearer 0.55 per cent once hard-to-learn tasks are discounted. Ten years of argument compressed into one spread: these are not competing estimates of the same object.

Acemoglu assumes — explicitly, and to his credit — that AI remains a tool: a cost saving applied to a fixed set of existing tasks, with no new tasks, no reorganisation of markets, no change in the process of science. Goldman's number quietly assumes the opposite: economy-wide reallocation, new occupations, compounding adoption. If agents merely assist workers at today's tasks, the sceptic's bound is roughly right and the whole episode is a rounding error on trend growth. If agents become **participants** — arbitraging prices, shopping for consumers, running research pipelines, underwriting loans, transacting with each other at machine speed — the gain arrives through channels Hulten's theorem was never asked to price, because the theorem holds task shares fixed while agents change what the tasks are.

FIG. 11.1 – THE MACRO RANGE IS A BET, NOT AN ERROR BAR

ACEMOGLU 0.55–0.71% TFP · 10yr

GOLDMAN +7% GDP · +1.5pp productivity



OECD and Aghion-Bunel sit toward the lower bound — on diffusion grounds, not micro-evidence [12]

The distance between 0.7 and 7 is the tool-versus-actor assumption. Sources: Goldman Sachs (Briggs & Kodnani 2023); Acemoglu (MIT 2024); OECD AI Paper 15; Aghion & Bunel (SF Fed 2024) [12].

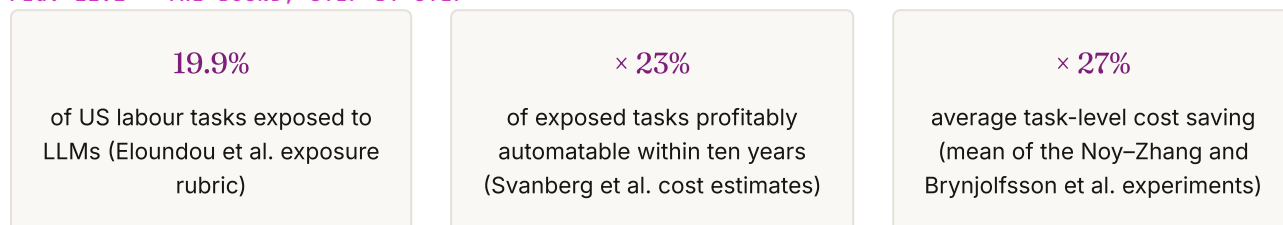
### THE BET, STATED

Same task data, same experiments, different ontology: is AI a cheaper way to do today's work, or a new class of economic actor? The next three pages show the sceptic's arithmetic in full, the field-evidence floor beneath both camps, and the indicators that will settle it.

## The sceptic's arithmetic, shown in full

An audit steelmans its strongest witness, and Acemoglu's calculation deserves to be displayed rather than summarised. By Hulten's theorem, aggregate TFP gains equal the GDP share of the tasks affected times the average cost saving on those tasks. Every input is drawn from the empirical literature this report has already cited — including the very experiments the optimists lean on:

FIG. 11.2 – THE BOUND, STEP BY STEP



Weighted into GDP via Hulten's theorem, the chain caps TFP at **≤0.71% cumulative over ten years** — about 0.07% a year; GDP effects ~1.1%, at most 1.6–1.8% with induced capital deepening. Discounting hard-to-learn tasks lowers the central bound to **0.55% TFP / 0.90% GDP**. Acemoglu, MIT 2024 [12].

Within its assumptions, the arithmetic is unimpeachable. The audit question is what sits outside the assumptions — because each exclusion is not an oversight but a deliberate scope line, and each one is a channel this report has separately priced:

- ▶ **New tasks.** The task set is held fixed at today's O\*NET catalogue; occupations and products that do not yet exist contribute nothing — yet new tasks are where past general-purpose technologies earned most of their return.
- ▶ **Reorganisation.** No reallocation across firms and markets, no new market structures, no Coasean redrawing of firm boundaries — the channel Analyses 2–4 priced in the trillions.
- ▶ **Science.** Acemoglu explicitly sets aside AI "changing the process of science" — the one channel with the Jones–Summers 5–13x multiplier standing behind it (Analysis 4).
- ▶ **Agents as actors.** No transacting, negotiating, arbitraging software participants at all — the entire subject of Movement III.


### ACEMOGLU'S DEEPER WARNING

His model also shows AI-enabled **new tasks can carry negative social value** — manipulation, additive design — such that AI could raise measured GDP ~2% while *reducing* welfare by 0.72%. The multiplier has no sign until institutions attach one. (MIT, Acemoglu 2024 [12])


If the assumptions hold through 2034, the bound is right and modest AI policy is correct policy. The empirical tripwires that would break them are specified on page 56 — several are already under measurement.

## The floor under both camps: what field evidence already shows

Beneath the forecast war sits a base that is no longer speculative: randomised and staggered deployments in real firms, an exposure census of the US task economy, and — since 2025 — live usage telemetry. Every number on this wall is measured, not modelled; and every deployed system studied so far was a suggestion tool, the **weakest possible form of agency**, which makes each figure a floor.



The last tile is the quiet hinge of the whole macro debate: the leap from 15 per cent of tasks to roughly half is precisely the leap **from model to agent**. And note the direction of the distributional finding — in both flagship studies AI assistance **compressed** skill gaps inside firms, the opposite of what three decades of skill-biased technical change taught economists to expect.



**The economy is in the transition zone: AI performs tasks at scale, mostly alongside humans, increasingly instead of them.**  
 Anthropic Economic Index, Handa, Tamkin et al. 2025 [13].

**FIG. 11.3 – USAGE IS FACT, NOT FORECAST**

Anthropic's Economic Index maps over four million Claude conversations onto the O\*NET task taxonomy: **~36% of occupations already use AI on at least a quarter of their tasks** (only ~4% on three-quarters or more). 57% of usage augments — learning, iterating, checking — while 43% automates, the model fulfilling the request with minimal human involvement.

## What would settle the bet

A disagreement this large, resting on named assumptions, is empirically decidable — and unusually for a macro dispute, the deciding instruments already exist. Acemoglu's bound fails only if its scope lines break: if tasks appear that O\*NET never catalogued, if automation displaces augmentation in the telemetry, if the cost of profitable automation keeps falling, or if agents begin transacting rather than suggesting. Each of those is observable quarter by quarter. The audit therefore closes this analysis not with a verdict but with an instrument panel — six indicators, each tied to a specific assumption, each with a named data source:

- 1 **The automation share.** Anthropic's telemetry currently reads 57/43 augmentation-to-automation. A sustained climb in the 43 — the model fulfilling requests, not assisting them — is the actor row arriving in the data. (Anthropic Economic Index [13])
- 2 **The profitable-automation fraction.** Svanberg's 23% is a 2024 cost estimate, and it falls as inference costs fall. Each revision mechanically loosens Acemoglu's bound — using his own arithmetic. (MIT 2024 [12])
- 3 **Tasks outside the catalogue.** The first appearance, at scale, of measured task portfolios containing work that did not exist in O\*NET breaks the fixed-task assumption directly. (EI telemetry vs O\*NET [13])
- 4 **The research-productivity number.** Nobody has yet measured agentic R&D's effect on research productivity at economy level. Because the Jones–Summers 5–13x multiplier stands behind this channel, that paper will be the most important growth-economics result of the decade. (Analysis 4; [04][12])
- 5 **Sectoral price dispersion.** If consumer and firm agents act as market participants, dispersion collapses Kerala-style, sector by sector; where it persists, the residual maps protected rents. (Jensen QJE 2007 template [15])
- 6 **The markup series.** De Loecker–Eeckhout's 18%→67% curve is the capture gauge: agentic commerce bending it down signals open competition; accelerating it signals the toll booth. This one decides the split, not the size. (NBER [06])

“

*The macro debate is really a debate about whether AI stays a tool or becomes an economic actor.*

ENSI · THE AGENTIC MULTIPLIER — REPORT 3 OF 3, 2026

What the bet is not about deserves one line: nobody serious disputes the field-evidence floor. Double-digit gains in real deployments, skill compression inside firms, half the task economy in scope once tooling counts — that is common ground. The dispute is confined to whether the machine's **mechanisms themselves** get upgraded. Which is why the audit's final analysis stops forecasting a number and starts mapping the futures the bet can land in.

MOVEMENT I · THE AUDIT – ANALYSIS 12

# 12 · Scenarios 2035

Eleven analyses reduce to two live uncertainties: does AI stay a tool or become an actor, and does the model layer stay open or concentrate? Crossed, they yield four futures for the value engine — and only one of the two axes is decided by technology.

Everything else in this audit is settled enough to carry forward. The mechanisms and their coefficients (Analyses 1–6), the debits and their repairs (7–8), the distribution record (9), the systems counterfactual (10) — these are the fixed physics of the scenario space. What Analysis 11 could not settle is the capability question: whether agents remain suggestion tools, in which case Acemoglu's bound holds and the decade adds a rounding error, or become transacting participants, in which case the field evidence and the amplification arithmetic point to an addition measured in trillions per year. That is the **row** of the matrix, and it will be decided mostly by laboratories and adoption curves.

The second uncertainty is older than AI. Nordhaus's 2.2 per cent — the audit's founding fact — was never a law of nature; it was produced by competition, imitation and patent expiry relentlessly passing surplus to consumers. Whether the agentic economy repeats that split depends on the market structure of its choke point: the model and agent layer. Open — interoperable, contestable, agent-access guaranteed — and the historical split holds. Concentrated — a handful of gatekeepers owning the agents that shop, negotiate and underwrite — and the 98/2 division can flip for the first time in the modern record. That is the **column** of the matrix, and it is decided by policy: protocol rules, access rights, competition enforcement. The one lever states actually hold.

| Row                                                                                                 | Column                                                                                                   | Stakes                                                                    |
|-----------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------|
| Capability — AI as tool vs actor; set by technology and adoption; watched via the Analysis-11 panel | Market structure — model layer open vs concentrated; set by policy; watched via markups and access rules | +0.5–1% TFP at the floor · +\$10–20T/yr at the ceiling · and who keeps it |

The four cells that follow are not predictions. They are the audited coefficients of this report run forward under each pair of assumptions — Acemoglu's bound for the tool row, the amplification arithmetic of CALCULATIONS §8 for the actor row, the Nordhaus split and the De Loecker–Eeckhout markup drift for the two columns.

**LABELLING RULE**

All scenario quantities on the following pages are **ENSI estimates** (back-of-envelope, labelled), assembled from sourced coefficients: Acemoglu  $\leq 0.71\%$  TFP [12] · agentic dividend \$10–20T/yr (CALC §8) · Nordhaus 2.2% split [02] · markup series 18%→67% [06] · Kerala compression [15].

The scenario space

One picture carries the whole analysis. The vertical axis is the Analysis-11 bet; the horizontal axis is the Nordhaus question. Today's telemetry places the economy just below the centre line and drifting upward — while the model layer, supplied by a handful of firms, starts the decade on the right-hand side.

FIG. 12.1 – FOUR FUTURES FOR THE VALUE ENGINE, 2035



Capability decides the row; openness decides the column — and only the column is a policy choice. Axes per Analysis 11 and the Nordhaus question; quadrant magnitudes are labelled ENSI estimates built from sourced coefficients (p57 labelling rule).

The colouring is deliberate: both right-hand cells are reachable from today's starting point without any technological surprise at all — concentration is the default commercial instinct of every platform. The left-hand column has to be built.

## The four futures, card by card

Each card carries the scenario's arithmetic, its surplus split, and the leading indicators that would signal its arrival. All magnitudes are **ENSI estimates** (BOE), assembled per the p57 labelling rule.

### S1 · The Rounding Error — tool × open

**GDP:** ~+0.5–1% TFP over the decade — Acemoglu's bound realised. **Split:** ~98/2 to society holds; gains widely shared, individually small.

**Leading indicators:** automation share plateaus near 43%; agent tooling adoption stalls at pilots; sectoral price-dispersion statistics static.

**ENSI estimate** · Acemoglu [12]

### S2 · The Quiet Toll — tool × concentrated

**GDP:** ~+0.5–1% TFP, same floor. **Split:** worsens — the 18%→67% markup drift of 1980–2014 extrapolates through the AI stack; small gains, privately gated.

**Leading indicators:** markups keep climbing; assistants default to walled gardens; API access tolls rise; no agent-access rights legislated.

**ENSI estimate** · De Loecker–Eeckhout [06] + [12]

### S3 · Kerala at Scale — actor × open

**GDP:** ~+\$10–20T/yr within the decade (CALC §8: research channel \$3–8T · price compression \$1–2.4T · misallocation repair \$4T+ · entry and trade effects). **Split:** consumers keep the bulk — the Nordhaus split survives its first technology transition.

**Leading indicators:** price dispersion collapses sector by sector; open agent-payment rails (x402-class) carry real volume; entry rates and young-firm skewness recover.

**ENSI estimate** · CALC §8 · Jensen [15] · Nordhaus [02]

### S4 · The Rent Engine — actor × concentrated

**GDP:** the same ~\$10–20T/yr is created — then tolled at the model gate. **Split:** capture flips; an economy that pays a private tax on its prices, competition, research and trade simultaneously — **the most efficient rent collector ever deployed.**

**Leading indicators:** model-layer shares concentrate; walled gardens gatekeep agent commerce; markups accelerate past the 2014 trendline; agent identity and access rules written by platforms, not law.

**ENSI estimate** · CALC §8 · [06][13]

Note what the matrix does *not* contain: a cell where the technology fails to matter at all — the field-evidence floor of Analysis 11 already rules that out — and a cell where concentration produces the big dividend and shares it. Concentration does not destroy the surplus; it re-addresses it. That asymmetry is the whole policy argument of the next page.

#### READING THE CARDS

The row changes the size of the prize by an order of magnitude; the column changes its owner. A state cannot choose its row — capability arrives on the laboratories' schedule — but every state chooses its column.

The policy-determined diagonal

Scenario matrices usually end with a shrug: four futures, pick your prior. This one does not, because the two axes are not symmetrical. The **row** — tool or actor — is set by model capability and adoption, which no mid-sized state controls and which the Analysis-11 instrument panel merely observes. The **column** — open or concentrated — is set by protocol rules, access rights, payment rails and competition enforcement: ordinary legislative material, decided jurisdiction by jurisdiction, in the next few years, while the equilibrium is still unformed. The audit's history supplies the precedent. The 2.2 per cent producer share that anchors this entire report was itself manufactured — by antitrust, patent expiry, open science and low entry barriers relentlessly handing surplus to consumers. The twentieth century **chose its column**, and the choice, not the technology, is what made capitalism the donation machine of Analysis 2.

Hence the diagonal that matters runs from the Quiet Toll to Kerala at Scale — from small gains privately gated to large gains publicly kept — and every step along it is a policy act, not a forecast: an open-agent-access rule, interoperable payment rails, agent identity and liability law, model-layer concentration watched the way central banks watch systemic banks. If capability stalls, those acts cost little and still bend the markup curve. If capability arrives, they are the difference between S3 and S4 — between the largest consumer-surplus event since the container ship and the largest toll booth ever built. Asymmetric stakes, symmetric cost: that is what "no-regret" means in an audit.

| WATCHBOARD INDICATOR                                 | AXIS IT MOVES                | READ IT FROM                                         |
|------------------------------------------------------|------------------------------|------------------------------------------------------|
| Automation share of AI usage (now 43%)               | Row — capability             | Anthropic Economic Index telemetry [13]              |
| Profitably automatable fraction (Svanberg 23%, 2024) | Row — capability             | Task-level cost studies, revised yearly [12]         |
| Sectoral price-dispersion indices                    | Row — agents as participants | National market telemetry (Jensen template [15])     |
| Markup trend vs the 18%→67% series                   | Column — capture             | De Loecker–Eeckhout method, updated [06]             |
| Model-layer market shares; walled-garden defaults    | Column — structure           | Model & agent-platform market data [13]              |
| Agent-access & identity rules on statute books       | Column — policy              | Legislation trackers, EU/US/UK (Report 3 build list) |

“

*Nordhaus's 2.2 per cent is not a law of nature but an achievement of competition policy.*

ENSI · THE AGENTIC MULTIPLIER — REPORT 3 OF 3, 2026

That sentence closes the audit. Movement I established what the machine pays, to whom, at what cost, against which alternatives, and under which futures. Movement II now opens the machine itself — the seven mechanisms that generate the payment, ranked by the value each creates.

# Seven mechanisms, ranked by the value they create

Capitalism is not one thing. It is a stack of seven distinct value-creating mechanisms — each measured separately by four decades of empirical economics, each with its own quantified dividend and its own documented failure mode. Treat the stack as one object and you get a referendum. Treat it as seven mechanisms and you get an engineering conversation: which components create the most value, which are degrading, which are cheap to maintain and catastrophic to break.

---


|   |                                                       |                                                |
|---|-------------------------------------------------------|------------------------------------------------|
| 0 | Institutions — the operating system                   | ×7 long-run income · AJR 2001                  |
| 1 | Innovation & creative destruction — the dominant term | \$13.30 per \$1 R&D · Jones-Summers            |
| 2 | The price system — the information layer              | factor-2 TFP when it breaks · Hayek 1945       |
| 3 | Competition — the enforcement mechanism               | ≈\$28B/yr, one industry · Borenstein-Rose      |
| 4 | Capital allocation — the steering system              | +1pp growth per depth quartile · Levine        |
| 5 | Entrepreneurship — the entry-exit experiment          | 3% of jobs, ~20% of job creation · Haltiwanger |
| 6 | Trade — the machine, extended across borders          | \$2.1T/yr ≈ \$18,131 per US household · PIIE   |

## MOVEMENT II · THE MACHINE – OVERVIEW

## The machine, exploded

Before the deep-dives, the schematic. Five working modules run on one information bus, and the whole assembly runs on one operating system. Rank order is by the social value each layer creates — and the output pipe leaks, by design, to everyone.

FIG. C-1 · THE SEVEN-MECHANISM STACK



Rank = estimated social value created. Rank 0 is off the scale: it does no producing itself, but hosts and multiplies all six producers.

**The stack, not a monolith.** Institutions host everything; prices coordinate everything; innovation creates most of the new value — and ~98% of that value leaks past its operators to society. That leak is not a flaw. It is the design.

Each module has been measured **separately, repeatedly, and by the best empirical economics of the past four decades**. Innovation converts research effort into permanently higher living standards. The price system aggregates dispersed knowledge no planner could collect. Competition forces the gains through to consumers. Finance steers capital toward its best uses. Entry and exit run the experiment that discovers which uses those are. Trade extends the whole machine across borders. And institutions set the expected payoff to every action the other six reward.

The literature does not just say "markets work." It says **which parts of the market machine do how much of the work** — and that is what this movement extracts, one mechanism per section, in a fixed anatomy that makes the seven directly comparable.

Sources for this movement: **NBER** Jones-Summers · Nordhaus · Levine · Haltiwanger **MIT** AJR 2001

**UChicago** Hayek 1945

**Stanford** Hsieh-Klenow

**Fraser** EFW 2024

**PIIE** Hufbauer-Lu

## MOVEMENT II · THE MACHINE – OVERVIEW

## How the ranking works

A ranking exists to allocate scarce reform capacity. A state cannot defend everything at once; it needs to know what to protect first and what it can safely leave alone. Three criteria produce the order, applied in sequence.

| CRITERION                        | THE QUESTION IT ANSWERS                                                                                                                          | DECISIVE EXAMPLE                                                                              |
|----------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------|
| <b>1 · Measured social value</b> | What is the quantified dividend per unit of input — or the cost of the mechanism's absence — on the strongest identification available?          | \$13.30 social return per \$1 of R&D (Jones-Summers) puts innovation first.                   |
| <b>2 · Substitutability</b>      | If this mechanism fails, can any other do its work? The less substitutable, the higher the rank.                                                 | No survey and no planner replaces the price signal — Hayek's argument holds prices at rank 2. |
| <b>3 · Leak rate</b>             | How much of the value escapes the operator and reaches society? A mechanism that enriches only its operator is a business, not a public machine. | Innovators keep ~2.2% of the surplus they create (Nordhaus) — a 97.8% leak.                   |

## THE AXIS

The vertical axis of this movement is **social value created** — not profit, not GDP share, not political salience. Where the literature offers a per-dollar return, we use it. Where the mechanism's value is a precondition rather than a flow (prices, institutions), we read it in the negative: what output collapses when the mechanism is absent or distorted.

## QUALITY OF EVIDENCE

Every headline number carries a confidence grade. **Very High**: peer-reviewed and replicated, with causal identification (AJR's settler-mortality instrument; Jensen's Kerala rollout). **High**: flagship institutional measurement (Fraser EFW; CEA; CMA). **Medium**: single study or back-of-envelope aggregation — always labelled, in this report, as an **ENSI estimate**.

### How to read each entry

Every mechanism is dissected with the **identical anatomy**, so the seven sections read as a matrix, not a pile of essays:

IN SHORT → WHY IT RANKS HERE → WHAT IT CONTRIBUTES → FAILURE MODE → THE AGENTIC HOOK → PRINCIPLES → EXAMPLES

One recurring finding across all seven sections — the closest thing this literature has to a law — deserves stating before the list: **the value each mechanism creates does not stay with the people who operate it**. Innovators keep ~2 percent; venture investors fund knowledge that spills to rivals; competitive firms hand efficiency gains to customers. The machine is a surplus-leaking engine, and the leak is the point.

MOVEMENT II · THE MACHINE – OVERVIEW


# The ranking at a glance

The whole movement in one strip. A Depth-2 reader can stop here; the deep-dives that follow supply the identification, the failure modes, and the agentic hooks.

| RANK | MECHANISM          | HEADLINE                                                                                | NUMBER | EVIDENCE                                 | ANCHOR | CONFIDENCE |
|------|--------------------|-----------------------------------------------------------------------------------------|--------|------------------------------------------|--------|------------|
| 0    | Institutions       | ×7 long-run income from institutional quality; freest quartile out-earns least-free     | 7.6×   | AJR, AER 2001 · Fraser EFW 2024          |        | Very High  |
| 1    | Innovation         | \$13.30 social benefit per \$1 R&D; social IRR 67%; society keeps ~98%                  |        | Jones-Summers, NBER 2020 · Nordhaus 2004 |        | Very High  |
| 2    | Price system       | Distorted prices cost a factor-of-2 in manufacturing TFP (read in the negative)         |        | Hayek, AER 1945 · Hsieh-Klenow, QJE 2009 |        | Very High  |
| 3    | Competition        | Fares ~30% below the regulated formula ≈ \$28B/yr — one industry, one year              |        | Borenstein-Rose, NBER 2007 · CEA 2016    |        | High       |
| 4    | Capital allocation | +1pp annual growth per financial-depth quartile; VC dollar up to 10× more patent-potent |        | Levine, NBER 2004 · Kortum-Lerner 1998   |        | High       |
| 5    | Entrepreneurship   | Startups: 3% of employment, ~20% of gross job creation                                  |        | Haltiwanger-Jarmin-Miranda, NBER 2010    |        | Very High  |
| 6    | Trade              | \$2.1T/yr to US income ≈ \$18,131 per household                                         |        | Hufbauer-Lu, PIIE 2017 · ACR 2010        |        | High       |

## The couplings, previewed

Ranking the mechanisms separately understates them jointly, because the strongest results in the library are about the **interactions**: import competition raised European firms' innovation and productivity (rank 3 feeds rank 1); venture capital exists to fund the high-variance young firms that create the disproportionate jobs (rank 4 feeds rank 5); prices let every module act at a distance; and the Hsieh-Klenow misallocation numbers show three mechanisms failing as one coupled system when institutions override prices. The machine's components are load-bearing for each other — the full treatment is the Compounding Stack, p. 93.



The build sequence. Each layer presupposes the ones to its left — which is why reform programs that start mid-chain (finance without courts, trade without prices) keep failing.

Two mechanisms carry live warning lights, flagged in their sections and read together in Movement I's debit ledger: **competition** (US markups up from 18% to 67% over cost, 1980–2014) and **entrepreneurship** (the high-growth skewness premium down from 16% to 4% since 2000). Both faults sit in the machine's **enforcement and access layers** — the wear surfaces — not in its core function. Nowhere in the 106-document library is there evidence that innovation stopped returning multiples, that prices stopped aggregating information, or that entry stopped creating jobs.

## RANK 0 • THE OPERATING SYSTEM

## 0 • Institutions — the operating system

## IN SHORT

Institutions — property rights, contract enforcement, constraint on the powerful — do no allocating, no innovating, no trading themselves. What they do is set the **expected payoff to every action the other six mechanisms reward**. Where they fail, the machine inverts: talent flows into rent-seeking, capital flees, prices are replaced by connections.

×7

long-run income gain implied by moving Nigeria's institutional quality to Chile's

AJR • AER 2001

7.6<sub>x</sub>

GDP per capita, freest vs least-free quartile: \$52,877 vs \$6,968

FRASER EFW 2024

\$7,610

poorest-decile income in the freest quartile — vs \$1,010 in the least free

FRASER EFW 2024

80.5<sub>yrs</sub>

life expectancy in the freest quartile, against 64.9 in the least free

FRASER EFW 2024

## WHY IT RANKS HERE

The catalog lists institutions seventh in the conventional ordering — and mis-lists them. The case for a low rank is flow accounting: in a functioning economy, institutions generate no measurable annual increment; innovation does. The case for the top is counterfactual accounting: the causal effect sizes attached to institutional quality **dwarf every other number in this movement**, because institutions do not add a term to the growth equation — they multiply all of them.

The identification problem — which causes which? — was cracked by Acemoglu, Johnson and Robinson. Where European colonisers could settle, they built inclusive, property-protecting institutions; where they died of disease, they built extractive ones; and settler mortality centuries ago — which cannot plausibly affect GDP today except through the institutions it selected — strongly predicts current income. Rodrik, Subramanian and Trebbi's horse race for world income differences — titled "**Institutions Rule**" — confirms it: control for institutional quality, and trade and geography largely work *through* it.

## THE RESOLUTION: RANK ZERO, NOT RANK SEVEN


You do not credit the operating system with the application's output — and you do not run the application without it. Every value estimate in this movement carries an implicit asterisk: *assuming enforceable property rights and functioning courts*.

RANK 0 · INSTITUTIONS – CONTINUED

WHAT IT CONTRIBUTES

In the Fraser Institute's 2024 *Economic Freedom of the World* index, the most economically free quartile of countries averages **\$52,877 of GDP per capita against \$6,968** in the least free — roughly eight to one. The same table settles distribution: the income *share* of the poorest decile is essentially identical across all four quartiles, while its absolute income multiplies by seven. **Free economies do not distribute a smaller share to their poor — they multiply the pie the identical share is cut from.**

FIG. C-2 · THE ECONOMIC-FREEDOM GRADIENT



The poorest decile of free economies out-earns the national average of unfree ones (\$7,610 vs \$6,968). Bars to a common dollar scale. Source: Fraser Institute, EFW 2024, ch. 1.

80.5 vs 64.9

life expectancy, yrs — freest vs least-free quartile (Fraser 2024)

9×

infant mortality in the least-free places relative to the freest (Fraser 2024)

~2.3–2.8%

poorest-decile income share — invariant across all quartiles (Fraser 2024)

The channels are concrete, not mystical. North's formulation — institutions as the humanly devised rules of the game that structure economic interaction — cashes out, in Besley and Ghatak's development treatment, as three payoffs: **secure property makes investment rational** (nobody irrigates a field that can be taken), **enforceable contracts make strangers tradeable partners** (extending exchange beyond kin), and **constrained executives make long-horizon commitments credible** (so factories and pension funds can exist).

Institutions are also the purest public good in the stack: the beneficiaries of secure property are overwhelmingly people who own almost nothing — because what property rights protect, at the margin, is **the small holder against the strong taker**. The Fraser poorest-decile line is the receipt: roughly \$6,600 a year of additional income for the bottom 10 percent, from rules they did not write. One honesty note: Haggard and Tiede find the rule-of-law-to-growth mapping surest at the extremes — and the World Bank's B-READY benchmarking (2024) is the gauge this section hands to Movement IV.

## RANK 0 · INSTITUTIONS – CONTINUED

## FAILURE MODE

Extraction and drift. The catastrophic mode is the state itself becoming the predator — AJR's extractive equilibrium, which their evidence shows can be **stable for centuries**: the extractive institutions colonisers installed where they could not settle still bind incomes today. The insidious mode is slower: rule-making captured by incumbents, rules accreting into entry barriers, until the institutions layer starts feeding the competition and dynamism faults documented at ranks 3 and 5. Advanced economies do not lose their operating system by coup; they lose it **by capture, clause by clause**.

## THE ASYMMETRY THAT JUSTIFIES RANK ZERO

Institutions pay ~zero measurable dividend per year in a functioning economy — and expose the entire \$150–180T/yr counterfactual flow when they break. Zero yield, total exposure: that is the profile of a liability to insure, not an asset to optimise.

## Proof of concept: Botswana

At independence in 1966, Botswana was among the poorest places on earth. Over the following ~35 years it posted **the fastest growth in the world** — with diamonds, which have wrecked other economies as often as they have enriched them. The difference, in Acemoglu, Johnson and Robinson's account, was institutional: pre-colonial constraint traditions carried into a post-independence state that protected property, enforced contracts, and disciplined its own executive. Same resource endowment as its neighbours; different operating system; opposite outcome.

Source: AJR, "An African Success Story: Botswana" (NBER w9124). The mirror-image natural experiments — Korea's 38th parallel, the two Germanys — put same-people, same-geography gaps at 10:1 to 20:1.

## THE AGENTIC HOOK

Institutions are the mechanism agents strain first and upgrade last. Movement III (pp. 103–105) shows both sides: **ambient compliance** — agents that make rule-following nearly free, turning regulation from a fixed cost into a background service — against the **governance-gap asymmetry**: agentic actors consume institutional capacity (contracts, identity, liability, adjudication) faster than institutions adapt. The operating system is about to receive ten thousand times more system calls. Whether it answers them is the multiplier's binding constraint.

Evidence chain: **MIT** AJR, AER 2001

**NBER** Rodrik-Subramanian-Trebbi

**Stanford/JEP** North 1991

**LSE** Besley-Ghatak 2010

**UT Austin** Haggard-Tiede 2010

**World Bank** B-READY 2024

**Fraser** EFW 2024

## RANK 0 · INSTITUTIONS — CONTINUED

“

*Innovation is where you invest; institutions are what you insure — and insurance you buy before the fire.*

ENSI · THE MACHINE, REPORT 2 OF 3 (2026)

## PRINCIPLES

- 1 **Protect absolutely; improve marginally.** The counterfactual exposure is seven-fold while the annual dividend is near zero — so the correct posture is vigilance against breakage, not perpetual redesign.
- 2 **Guard the extremes first.** The evidence maps institutional quality to growth most reliably at its extremes (Haggard-Tiede): preventing descent matters more than fine-tuning the top.
- 3 **Watch capture, not collapse.** In mature democracies the OS degrades by incumbent-written rules accreting into entry barriers — audit the rulebook the way rank 3 audits markups.
- 4 **Measure continuously.** Governance indicators and B-READY-style benchmarking are the machine's oil-pressure gauges; a stress-test each political cycle, because capture tracks elections, not wear.
- 5 **Keep the poorest decile on the dashboard.** \$7,610 vs \$1,010 is the institution dividend's receipt — the constituency of good rules is precisely the people who own least.

## EXAMPLES

- ▶ **Botswana** — fastest ~35-year growth record in the world on pre-colonial constraint institutions plus property rights (AJR, NBER w9124).
- ▶ **Korea, twice** — one people, one peninsula, two operating systems: an order-of-magnitude income gap opened within two generations of the 1948 split.
- ▶ **Estonia** — rebuilt the OS first (1992 radical liberalisation: property, currency, courts), then ran the digital state on top of it (Laar).
- ▶ **Nigeria → Chile** — the AJR counterfactual: closing the institutional-quality gap implies up to a 7-fold long-run income gain, against an actual gap above 11-fold.
- ▶ **B-READY 2024** — the World Bank's successor benchmarking regime: institutions made measurable, comparable, and therefore maintainable.

**The close: insure, don't invest.** For a policymaker asking *where does my next unit of reform effort earn most*, institutions are usually the wrong answer in a functioning democracy — the 67% social return at rank 1 wins. For the same policymaker asking *what must I never allow to break*, institutions win and it is not close. Judicial independence, property security, constraint on expropriation: zero measured annual dividend, total counterfactual exposure. The premium on this insurance is vigilance, and it is the cheapest line in the national budget.

## RANK 1 · THE DOMINANT TERM

# 1 • Innovation & creative destruction

## IN SHORT

Every dollar invested in research returns roughly **thirteen dollars to society** — five under deliberately hostile assumptions, twenty-plus counting health gains. No other expenditure category in public or private life has a documented return in this range. And its proceeds go overwhelmingly to people who never invested a cent.

**\$13.30**

social benefit per \$1 of R&D, baseline economy-wide estimate

JONES-SUMMERS · NBER  
W27863

**67%**

social internal rate of return on R&D — vs 10–15% private returns

JONES-SUMMERS · NBER  
W27863

**\$4.90**

the return that survives a 20-year diffusion-delay stress test (IRR still 11%)

JONES-SUMMERS · NBER  
W27863

**2.2%**

share of the surplus innovators keep — society banks the other ~98%

NORDHAUS · NBER  
W10433

## WHY IT RANKS HERE

First on criterion one, and it is not close. The load-bearing number of this entire report comes from Jones and Summers, who built an economy-wide calculation that nets out every spillover margin at once — imitation, business stealing, intertemporal effects — by reading the social return directly off the path of GDP. Their baseline: **\$1 of R&D produces \$13.30 of benefit in present value**, a social internal rate of return of 67 percent. Stress-test it and the number bends but does not break: impose a 20-year diffusion delay consistent with the micro-evidence and the benefit is still \$4.90 per dollar. Their own deliberately conservative summary is about \$5 per \$1, with plausible adjustments for health gains, inflation mismeasurement and international spillovers pushing it to \$10, even \$20 — measured for the US economy alone.

Innovation also wins on substitutability: prices coordinate, competition enforces, finance steers — but only innovation **moves the frontier**. Every other mechanism in this movement is, in the end, plumbing for this one: prices signal where the improvement is worth making, competition forces incumbents to respond or die, finance funds the challenger, entrepreneurship carries the idea into a firm.

## THE STANDING ARBITRAGE


An asset class returning 67% a year would, in any private portfolio, be scaled up until the return was arbitrated away. Society has never done the scaling — total US R&D is ~2.7% of GDP. That gap is the single largest identified under-investment in the evidence base.

## RANK 1 · INNOVATION – CONTINUED

## The engine: how creative destruction works

Growth in the Schumpeterian account is not accumulation but **replacement**: quality-improving innovations that render the old obsolete, formalised in the Aghion-Howitt model, where the engine of rising living standards is precisely the destruction of incumbent rents by better products. The market's contribution is to make this replacement **routine** — to give any firm with a better idea both the incentive (a temporary monopoly profit) and the right (free entry) to displace the installed base. The temporary rent is the carrot; imitation, patent expiry and competitive entry are the scheduled confiscation; and the confiscated surplus — lower prices, better products, freed ideas — is what society collects.

FIG. C-3 · THE CREATIVE-DESTRUCTION LOOP



The loop runs continuously; each pass raises the quality frontier and hands ~98% of the increment to people outside the innovating firm.

**Replacement, not accumulation.** Aghion-Howitt formalise growth as a loop in which the reward for innovating is designed to be temporary — and the leak to society is the machine's output, not its waste. Sources: Aghion & Howitt (NBER); Nordhaus (NBER w10433).

Hold the input in view alongside the return, because the ratio is the policy argument. The denominator of the Jones-Summers calculation — total US R&D across business, government, universities and non-profits — is only about **2.7 percent of GDP**; venture capital, the sharpest allocation channel at rank 4, reached roughly \$130 billion in 2018 but has often run below \$30 billion a year. The entire innovative effort of the largest economy on earth is a rounding error against its own output, earning a return no pension fund will ever see.

RANK 1 · INNOVATION – CONTINUED

WHAT IT CONTRIBUTES

The Jones-Summers estimate is deliberately a range, and every point in the range clears every hurdle rate in use anywhere. The conservative bound imposes long diffusion lags and strips out everything hard to measure; the upper bound adds what is real but usually uncounted — **health gains, quality mismeasurement, and spillovers across borders**. The firm-level evidence agrees from below: Bloom, Schankerman and Van Reenen separated R&D's two countervailing spillovers — positive knowledge spillovers to technological neighbours versus negative business-stealing effects on product-market rivals — and found the technology spillovers **quantitatively dominant**, so social returns run about twice private returns and the economy systematically under-invests. Nor is the return confined to private labs: long-run productivity returns to *government* R&D are also high (Dallas Fed, Fieldhouse-Mertens 2023).

FIG. C-4 · SOCIAL BENEFIT PER \$1 OF R&D



**Even the hostile scenario returns 5-to-1.** All variants from the same economy-wide model; the baseline corresponds to a 67% social IRR. Source: Jones & Summers, NBER w27863.

|                                   |                                                     |                                                                       |                                               |
|-----------------------------------|-----------------------------------------------------|-----------------------------------------------------------------------|-----------------------------------------------|
| 67%                               | 10–15%                                              | 2×                                                                    | 2.7%                                          |
| social IRR on R&D (Jones-Summers) | typical private return — the wedge is the spillover | social vs private returns from technology spillovers alone (BSV 2007) | of GDP — total US R&D, the under-scaled input |

Scale the ratio to the world and the magnitude shows: global R&D spending is roughly \$2.9 trillion a year (2025). At \$5–13.30 of social value per dollar, the world's research effort generates on the order of **\$14.5–38.6 trillion a year of social value** — of which private actors capture a small fraction. The capitalist innovation system is society's highest-return asset class, and it is systematically *under-supplied* — which is exactly why the 67 percent social return persists instead of being arbitrated away.


ENSI ESTIMATE

The \$14.5–38.6T/yr global aggregation is a back-of-envelope extension of the Jones-Summers US ratios to world R&D spend — labelled accordingly. The per-dollar ratios themselves are peer-reviewed, US-measured, and conservative by construction.

RANK 1 · INNOVATION – CONTINUED

The giveaway: why 98 percent escapes

Because the innovator cannot keep it. Nordhaus, measuring "Schumpeterian profits" across the US non-farm business economy from 1948 to 2001, found that producers captured about **2.2 percent** of the total social surplus generated by technological advance (range 0.8–3.3%; such profits were just 3.8% of corporate profits). The rest passed to consumers through lower prices and to later innovators through the ideas themselves. The mechanism most associated in the public mind with billionaire fortunes is, in measured fact, a machine that gives away roughly 98 percent of what it creates. **The fortunes are real; they are the 2 percent.** Inverted, 2.2% is a ~45× giveaway multiplier (ENSI estimate): every \$1 billion of genuine innovation profit marks roughly \$44 billion of surplus that went to everyone else.



FAILURE MODE

Two faults, and they compound. First, **under-investment by design**: if society keeps 98 percent of the value, the private investor sees only a sliver of the return and rationally spends too little — which is why the social return sits at 67 percent instead of being competed down. Second, **the ideas are getting more expensive**: research productivity has been falling for decades across fields. It now takes more than 18 times the researchers of the early 1970s to sustain the Moore's-Law doubling, and aggregate US research productivity has declined by a factor of 41 since the 1930s — roughly –5.1 percent a year (Bloom, Jones, Van Reenen & Webb, 2020). Neither fault indicts the mechanism; both say the same thing: the binding constraint on the machine's dominant term is the **volume of innovative effort**, and the fix earns a measured 30–70 percent social return.

18×

more researchers needed today to hold the Moore's-Law cadence vs early 1970s

÷41

aggregate US research productivity since the 1930s

–5.1%/yr

the average rate of research-productivity decline

~45×

the giveaway multiplier implied by 2.2% capture (ENSI estimate)

Ideas are getting harder to find at exactly the moment a new class of researcher becomes available — the setup for this mechanism's agentic hook, overleaf.

## RANK 1 · INNOVATION — CONTINUED

## DIRECTION, NOT JUST VOLUME

What gets invented follows incentives as surely as how much. WIPO's direction-of-innovation work documents that the composition of world invention tracks demand, relative prices and policy signals — innovation is steerable, which makes the price system (rank 2) and competition (rank 3) part of research policy. The OECD's STI Outlook adds the fiscal warning: public research systems are under strain in most member states even as the measured social return argues for expansion. The machine's best component is simultaneously under-funded and mis-aimed — two dials, both adjustable.

## THE AGENTIC HOOK

Innovation's binding constraint is the volume and cost of research effort — precisely the input AI agents attack. Movement III (pp. 99–102) runs the arithmetic: research agents that cut effective research cost by 20–30 percent act as a permanent enlargement of the effective R&D base, and the Jones-Summers multiplier converts that into **+\$3–8 trillion a year of social value (ENSI estimate)**. The 18×-researchers problem is the first constraint in this movement with a plausible technological answer — and because of the 2.2% split, ~98% of whatever agents add here is pre-committed to society, so long as competition survives the transition.

## PRINCIPLES

- 1 **Fund the volume.** The under-investment gap is the standing arbitrage: every marginal R&D dollar — public or private — earns a measured 30–70% social return until the gap closes.
- 2 **Subsidise R&D, not smallness.** Knowledge spillovers run through large firms at least as strongly as small ones (BSV 2007) — target the activity, not the company size.
- 3 **Protect the destruction half.** The loop stalls if incumbents can buy the rules: creative destruction is a competition-policy outcome, not just a science-policy one.
- 4 **Count health and diffusion.** The true return is \$20+ per dollar once health gains and international spillovers are counted — budget against that number, not the patent ledger.
- 5 **Expect no gratitude arithmetic.** A 2.2% capture rate means the winners will always look overpaid while society banks 45× their take — read fortunes as receipts, not costs.

## EXAMPLES

- ▶ **Moore's Law** — held for five decades only by scaling researcher headcount 18× — the clearest picture of ideas getting harder to find (Bloom-Jones-Van Reenen-Webb).
- ▶ **The smartphone category** — Apple's ~\$1T cumulative net income against consumer surplus measured in the tens of trillions at the Nordhaus split (ENSI estimate).
- ▶ **Health R&D** — the margin that lifts the multiplier past \$20: mortality gains are the largest uncounted line in the innovation ledger (Jones-Summers).
- ▶ **Government R&D** — long-run productivity returns comparable in kind to private R&D: the public half of the engine is not a subsidy but an investment (Dallas Fed 2023).
- ▶ **The VC-backed cohort** — 89% of R&D among surviving IPO-era firms came from venture-funded companies: rank 4's channel feeding rank 1's engine (Lerner-Nanda).

Evidence chain: **NBER** w27863 Jones-Summers **NBER** w10433 Nordhaus **NBER** BSV 2007

**Stanford/AER** Bloom et al. 2020

**WIPO** direction of innovation

**OECD** STI Outlook

**Dallas Fed** Fieldhouse-Mertens

## RANK 2 · THE INFORMATION LAYER

## 2 · The price system

## IN SHORT

The economic problem is not allocation given full information — it is that the relevant knowledge exists only as **dispersed, contradictory, local fragments** in millions of separate minds, fleeting and tacit, collectable by no statistical bureau. Prices solve it: a single number that summarises everything anyone needs to know to act, published continuously, for free, to everyone.

## WHY IT RANKS HERE

Second by a peculiarity of accounting: the price system's dividend cannot be measured per dollar the way R&D can, because it is **the precondition for the accounting itself**. Its value shows up only as the difference between system-level outcomes when it is present versus absent — a difference the twentieth century measured at continental scale, and which the misallocation literature (p. 76) prices at a factor of two in manufacturing productivity. On the substitutability criterion it is untouchable: no survey, no census, no planner has ever replicated what the signal does, and Hayek's argument is that none can even in principle — the facts the market coordinates do not exist before the market process generates them.

Hayek's 1945 demonstration remains the deepest single insight in the library, and it turns on one commonplace example:

“

*Assume that somewhere in the world a new opportunity for the use of some raw material, say tin, has arisen, or that one of the sources of supply of tin has been eliminated. It does not matter for our purpose — and it is very significant that it does not matter — which of these two causes has made tin more scarce. All that the users of tin need to know is that some of the tin they used to consume is now more profitably employed elsewhere, and that in consequence they must economize tin.*

**F. A. HAYEK** · "THE USE OF KNOWLEDGE IN SOCIETY," AMERICAN ECONOMIC REVIEW, 1945, P. 526  
— VERIFIED AGAINST THE LIBRARY COPY

The users need never learn *why*. The price rises; every user economises; every producer of substitutes expands; the substitutes of the substitutes adjust; and the effect propagates through the whole economic system — **without any participant knowing the cause and without any central instruction**. The whole acts as one market not because anyone surveys the field, but because individual fields of vision overlap enough for the relevant information to reach all, through a number.

## RANK 2 · PRICES — CONTINUED

## The marvel of unplanned coordination

Hayek's description of the machinery is literal, not poetic: the price system is "**a system of telecommunications**" in which the producer watches "the movement of a few pointers, as an engineer might watch the hands of a few dials," adjusting to changes of which he may never know more than is reflected in the price movement. In abbreviated form, by a kind of symbol, only the most essential information is passed on — and only to those concerned. The economy of knowledge with which the system operates is its most significant property: how little each participant needs to know in order to take the right action.

“

*The marvel is that in a case like that of a scarcity of one raw material, without an order being issued, without more than perhaps a handful of people knowing the cause, tens of thousands of people whose identity could not be ascertained by months of investigation, are made to use the material or its products more sparingly; i.e., they move in the right direction.*

F. A. HAYEK · AER 1945, P. 527 — VERIFIED AGAINST THE LIBRARY COPY

Hayek chose the word deliberately — "to shock the reader out of the complacency with which we often take the working of this mechanism for granted." Had the price system been the product of deliberate human design, he wrote, it "would have been acclaimed as one of the greatest triumphs of the human mind." Its misfortune is double: it is not the product of design, and the people guided by it usually do not know why they are made to do what they do. His companion essay completes the argument: competition is a **discovery procedure** — the facts the market coordinates (who can produce what, at what cost, for whom) are not data waiting to be collected but outputs of the market process itself, which is why no questionnaire can substitute for it.

### WHAT IT CONTRIBUTES

Nobody owns the price system. Its entire output — the continuous, free publication of what everything is worth to everyone else — is **a public good produced as a by-product of private transactions**. Every participant free-rides on the information generated by all the others, and there is no toll booth at which an owner could stand. On this movement's third criterion, the leak rate, prices are the limiting case: a mechanism whose operators capture nothing at all, because there are no operators — only users.

#### HAYEK'S CAVEAT, KEPT HONEST

The system fulfills its function "less perfectly as prices grow more rigid" — administered, subsidised or frozen prices are degraded bandwidth on the machine's only bus. That is the bridge to the negative proof overleaf.

## RANK 2 · PRICES — CONTINUED

## The negative proof: what broken prices cost

Because the price system's value is a precondition, the cleanest measurements are of its **absence**. The comparative literature on state versus private ownership finds private, price-disciplined operation systematically outperforming state operation across industries and countries, for information and incentive reasons rather than ideology (Shleifer 1998; Djankov et al. 2003).

The marginal estimate comes from within market economies: Hsieh and Klenow measured marginal products of capital and labour across manufacturing plants *within narrow industries* in China and India — where subsidised credit, administered inputs and connections override prices — and found dispersion a working price system would arbitrage away:

FIG. C-5 · PRICE DISTORTION, MEASURED TWO WAYS



Where prices are overridden, identical-industry plants diverge by a factor of 6–7 — and repairing the signal is worth up to a doubling of manufacturing productivity. Source: Hsieh & Klenow, QJE 2009.

Misallocation is not a rounding error; it is a large fraction of why poor countries are poor — and it is **the price system's value read in the negative**. China's own within-sample improvement — reallocation gains of roughly 2 percent of TFP per year, 1998–2004 — tracks the market reforms Movement I credited.

## FAILURE MODE

Prices transmit the knowledge people act on; they fail where information is asymmetric or where costs fall on parties outside the transaction. Stiglitz's formalisation is the honest steelman: with imperfect information and incomplete markets, the invisible-hand theorems do not hold exactly, and market outcomes are not constrained-efficient. The engineering response is bounded: **targeted instruments** — disclosure rules, insurance design, externality pricing — layered on the signalling layer, never a replacement for it, since the planner's information problem is strictly worse.

## THE LARGEST OPEN TICKET


Carbon remains the biggest unpriced externality in the system — the one place where the bus carries a wrong number by design. Pricing it is maintenance of rank 2, not a departure from it.

## RANK 2 · PRICES — CONTINUED

## Preview: when the information layer switches on

The cleanest modern experiment in this movement: as mobile-phone coverage rolled out region by region along the coast of Kerala, fishermen at sea could, for the first time, learn prices at every beach market before choosing where to land. Price dispersion across markets — huge and persistent for as long as records exist — **collapsed to the law of one price within days** of each region's towers switching on. Waste, previously 5–8 percent of the catch dumped unsold, fell to **near zero**; fishermen's profits rose **8 percent**; consumer prices fell **4 percent**. Better information created value for both sides at once — nothing was redistributed, and nothing new was produced except coordination.

FIG. C-6 · THE KERALA PRICE CONVERGENCE (SCHEMATIC)



One phone signal did what no regulator could: it turned three noisy local markets into one informed one. Schematic reconstruction of the pattern in Jensen, QJE 2007 — the sardine study; series stylised, magnitudes as published.

## THE AGENTIC HOOK

Kerala is the price system's upgrade demo — and it ran on 2G voice calls. Movement III (pp. 99–102) makes the extension: AI agents that query, compare and transact continuously are **Kerala run economy-wide and permanently** — compressing retail and services price dispersion across the ~\$60T global consumption base for an estimated \$1–2.4T a year of surplus (ENSI estimate), and turning residual price dispersion into a live map of where rents still hide.

## PRINCIPLES

- 1 **Leave the layer alone.** The price system needs no subsidy and tolerates no substitute — the only mechanism in the stack for which non-interference is the whole policy.
- 2 **Patch failures with targeted instruments.** Disclosure, insurance design, externality pricing — corrections ride on the signal; they must never replace it.
- 3 **Finish pricing the externalities.** Carbon is the layer's one open ticket: a wrong price is worse than a missing one.
- 4 **Treat dispersion as diagnosis.** Where identical goods hold different prices, information or competition is failing — Hsieh-Klenow at the plant level, Jensen at the beach.
- 5 **Fear rigidity, not volatility.** Administered and frozen prices degrade the bandwidth of the economy's only bus (Hayek's own caveat).

## EXAMPLES

- ▶ **Hayek's tin** — one number moves thousands of unknowing actors in the right direction, worldwide (AER 1945, verified above).
- ▶ **Kerala fisheries** — dispersion → law of one price within days of coverage; waste ~0, profits +8%, consumer prices -4% (Jensen, QJE 2007).
- ▶ **China & India misallocation** — overridden prices cost a factor-of-2 in manufacturing TFP (Hsieh-Klenow, QJE 2009).
- ▶ **State vs private ownership** — the A/B test of administered versus price-guided production, decided on information grounds (Shleifer 1998; Djankov et al. 2003).
- ▶ **Polycentric commons** — Ostrom's documented class of price-failure corrections that are neither market nor state (Nobel lecture, 2009).

Evidence chain: **UChicago/AER** Hayek 1945 **Mises reprint** Hayek, Discovery Procedure **QJE** Hsieh-Klenow 2009

**QJE** Jensen 2007

**NBER** Shleifer 1998 · Stiglitz 1991

**Nobel** Ostrom 2009

## MOVEMENT II · THE MACHINE · MECHANISM DEEP-DIVE

### 3 Competition — the enforcement mechanism

## IN SHORT

Innovation creates the surplus and prices signal where it lies; competition is what forces the surplus **through** to consumers. A firm facing rivals cannot hold prices above cost for long, or let quality slide into the quiet life of monopoly. Competition is the machine's enforcement layer: it converts the **capacity** to deliver value into the **obligation** to deliver it.

## -30%

US airline fares vs the regulation-era SIFL formula, by 2005 — deregulation as natural experiment

NBER · Borenstein & Rose 2007

## \$28B/yr

Consumer welfare gain from that one deregulated industry, in that single year

NBER · Borenstein & Rose 2007

## 18→67%

Average US markup over marginal cost, 1980 → 2014 — the warning light on this mechanism

NBER · De Loecker & Eeckhout 2017

## >5<sub>x</sub>

90th-percentile firm's return on invested capital vs the median — was ~2x a quarter-century ago

White House CEA 2016

## WHY IT RANKS HERE

Competition ranks third because it is the mechanism that makes the top two **pay out**. Every value estimate in ranks 1 and 2 quietly assumes the gains reach society; competition is that assumption made real. The 2.2 per cent producer share documented in Movement I is competition's arithmetic signature — it is rivalry, imitation and entry that strip roughly 98 per cent of the innovation surplus from the innovator and hand it to everyone else (NBER, Nordhaus, 2004).

It also passes the substitutability test that orders this catalog: **nothing else disciplines**. A regulator can cap one price; only rivalry polices price, quality, cost, variety and management simultaneously, continuously, and without an administrator. The White House Council of Economic Advisers' synthesis states the mechanism plainly: competition lowers prices, raises quality and variety, drives innovation and productivity, and reaches workers too, because firms must compete for labour as well as customers (CEA, 2016). Competition ranks below innovation and prices only because it **moves** surplus rather than creating it — but the engine is socially worthless without its enforcement layer.

Evidence base for this mechanism:

**NBER** Borenstein & Rose 2007

**LSE CEP** Bloom & Van Reenen 2006

**LSE CEP** Bloom, Draca & Van Reenen 2011

**CEA** 2016

**CMA** 2024

**OECD** 2014

**NBER** De Loecker & Eeckhout 2017

WHAT IT CONTRIBUTES

Three independent measurement strategies converge on the same answer — the triangulation, not any single number, earns the Very High confidence grade. And because competition's very definition is the transfer of surplus from producers to consumers, the leak rate here approaches its maximum by construction.

**Deregulation as natural experiment.** The 1978 Airline Deregulation Act removed federal price and entry controls from one large, well-documented industry — and left behind the counterfactual in writing, because the old regulatory fare formula (SIFL) continued to be published. By 2005, actual US fares ran about **30 per cent below** what the regulation-era formula would have charged, a consumer welfare gain of **\$28 billion in that single year, in that single industry** (NBER, Borenstein & Rose, 2007). Scaled across trucking, telecoms, rail and energy, the deregulation era's consumer dividend runs to hundreds of billions.

**The management channel.** Competition disciplines competence, not just prices. Bloom and Van Reenen's management-practice

surveys, scored across thousands of firms and many countries, found that better management robustly predicts higher productivity, profitability and survival — and that **weak product-market competition is one of the two main explanations for the long tail of badly managed firms**, accounting, together with family succession, for about half of it (LSE CEP, Bloom & Van Reenen, 2006). Management quality is among the strongest known drivers of TFP differences between firms — a productivity mechanism in its own right.

**Rivalry as a technology-forcing device.** When Chinese import competition intensified in European product markets, the affected European firms did not merely shrink — they **increased patenting, IT adoption and productivity** (LSE CEP, Bloom, Draca & Van Reenen, 2011). Competition does not only redistribute a fixed surplus; it forces the creation of new surplus, feeding rank 1 directly — and the OECD's cross-country evidence closes the loop, linking competition-policy strength to productivity, employment and lower prices (OECD, 2014; CMA, 2024).

| STRATEGY                 | DESIGN                                                     | HEADLINE RESULT                                                                                 | SOURCE                                  |
|--------------------------|------------------------------------------------------------|-------------------------------------------------------------------------------------------------|-----------------------------------------|
| Deregulation experiment  | Post-1978 fares vs published SIFL counterfactual formula   | Fares <b>-30%</b> by 2005; ≈ \$28B consumer gain in one year, one industry                      | NBER, Borenstein & Rose 2007            |
| Management surveys       | Scored practices across thousands of firms, many countries | Weak competition explains much of the <b>badly-managed tail</b> (with family succession, ~half) | LSE CEP, Bloom & Van Reenen 2006        |
| Trade-induced innovation | Chinese import exposure across European firms              | Patenting, IT adoption and TFP <b>rise</b> under import rivalry                                 | LSE CEP, Bloom, Draca & Van Reenen 2011 |
| Official syntheses       | Economy-wide indicator reviews, US & UK & OECD             | Competition → lower prices, higher quality, wages and productivity                              | CEA 2016 · CMA 2024 · OECD 2014         |

## FAILURE MODE

Market power — and here the warning light is genuinely on. De Loecker and Eeckhout, measuring markups across US publicly traded firms, document a rise in the average markup from **18 per cent above marginal cost in 1980 to 67 per cent in 2014** — a three-and-a-half-fold increase, concentrated in the upper tail of the firm distribution rather than spread evenly (NBER, De Loecker & Eeckhout, 2017). The CEA's indicators point the same way: revenue shares of the top 50 firms rose in most broad sectors between 1997 and 2012 — transportation and warehousing by 11.4 percentage points, retail by 11.2 — and the 90th-percentile firm's return on invested capital now runs at **more than five times the median**, a ratio that stood near two a quarter-century earlier (CEA, 2016). The UK's economy-wide stocktake finds parallel, milder trends in concentration and markups (CMA, 2024).

FIG. C3-1 — THE MARKUP DRIFT



The average US markup over marginal cost roughly tripled in 34 years — every point of markup is a point of surplus that stops flowing to consumers. Source: NBER, De Loecker & Eeckhout 2017 (US publicly traded firms).

FIG. C3-2 — WINNERS PULLING AWAY



Top-firm returns on invested capital have detached from the median — persistent, outsized returns are the classic signature of weakened rivalry. Source: White House CEA 2016.

## WHY THIS FAULT MATTERS MOST

Market power does not just tax one market — it corrodes the **leak rate of the whole machine**, the very property that makes capitalism socially valuable. The same fault also feeds rank 5's dynamism decline downstream.

The engineering read: this is the **best-understood fault in the machine**, with a two-century-old toolkit — antitrust enforcement, merger control, entry-barrier removal, occupational-licensing reform. The markup series is best understood as a maintenance backlog, not a design flaw. A machine with a worn brake pad needs a brake job, not a referendum on braking.

## THE AGENTIC HOOK

A large share of modern market power lives not in patents or scale but in **friction**: consumers who cannot compare, switching costs deliberately inflated, drip pricing, fee schedules written to defeat attention. An assistant agent that carries your context, interrogates every rival's offer, reads every footnote and re-runs the comparison monthly reduces search costs to approximately zero; obfuscation pricing does not survive a counterparty with infinite patience (Microsoft Research, Rothschild et al., 2025). The conditional is architectural: the same analysis distinguishes an open "web of agents", in which rivalry tightens toward the Bertrand textbook, from **agentic walled gardens**, in which the shopping agent is owned by the platform being shopped — the search-cost moat nationalised by the gatekeeper. Amplification arrives in both worlds; **the question is for whom**. Protocol policy — agent access to offers, interoperability, portable data — is the competition policy of the 2030s.

## PRINCIPLES

- 1 **Treat markups as telemetry.** Publish national markup, concentration and entry gauges annually — the De Loecker–Eeckhout series is a dashboard instrument.
- 2 **Police mergers at the entry margin.** The costliest harm is the challenger never born; merger review should weigh future rivalry, not only present prices.
- 3 **Demolish entry barriers before subsidising anyone.** Licensing reviews and standardised market access return surplus without spending a cent of fiscal capacity (CEA, 2016).
- 4 **Keep labour markets rivalrous too.** Non-compete limits and wage-fixing enforcement extend the mechanism to the input side (CEA, 2016).
- 5 **Extend competition law to protocols.** A consumer agent's right to read any published offer on non-discriminatory terms is the decisive instrument of the agentic decade.

## EXAMPLES

- ▶ **US airlines, 1978–2005** — deregulated fares settle ~30% below the surviving regulatory formula; ≈ \$28B/yr consumer gain in one industry (NBER, Borenstein & Rose 2007).
- ▶ **European firms under Chinese import competition** — patenting, IT adoption and productivity rise under rivalry pressure (LSE CEP, Bloom, Draca & Van Reenen 2011).
- ▶ **Cross-country management surveys** — weak competition (with family succession) explains about half of the badly-managed tail of firms (LSE CEP, Bloom & Van Reenen 2006).
- ▶ **United Kingdom, 2024 stocktake** — concentration and markup trends parallel the US in direction, milder in size (CMA, State of UK Competition 2024).

## MOVEMENT II • THE MACHINE • MECHANISM DEEP-DIVE

## 4 Capital allocation — the steering system

## IN SHORT

An economy's growth depends less on how much it saves than on **where the savings go**. Financial systems earn their keep by doing five things no household can do alone: producing information about investments, monitoring managers, spreading risk, pooling savings, and easing exchange (NBER, Levine, 2004). When these work, capital flows to its highest-return uses.

### +1<sub>pp/yr</sub>

Extra per-capita growth from moving financial depth from the slowest- to the fastest-growing country quartile

NBER • Levine 2004  
(King-Levine)

### 3–10<sub>x</sub>

Patent potency of a venture-capital dollar vs a corporate-R&D dollar

NBER • Kortum & Lerner 1998

### 47%

Of all 4,109 US IPOs 1995–2018 were venture-backed — from a channel reaching <0.5% of new firms

HBS • Lerner & Nanda 2020

### ×2<sub>TFP</sub>

Manufacturing productivity forgone to capital misallocation in China and India — the cost of this mechanism's absence

QJE • Hsieh & Klenow 2009

## WHY IT RANKS HERE

Capital allocation ranks fourth on the strength of a rare double proof: the mechanism has been measured **working** — cross-country, instrumented, and at the industry level — and measured **absent**, in the misallocation accounting. Both estimates agree on magnitude: roughly a doubling of manufacturing productivity separates well-steered from politically steered capital.

It ranks below competition for one reason: partial substitutability. Firms can and do finance investment from retained earnings; an economy with weak intermediation limps rather than halts. But the substitution is bitterly regressive — internal finance favours incumbents with cash flows over challengers with ideas. The steering system is rank 5's and rank 1's fuel line.

## “

*When the steering fails, capital pools around collateral, connections and the state — and the economy runs its best projects on empty tanks.*

ENSI • THE MACHINE — REPORT 2 OF 3, 2026

## WHAT IT CONTRIBUTES

The evidence stacks in four layers, from macro correlation to micro accounting — each layer answering the previous layer's objection.

**Layer one: the cross-country gradient.** In the King–Levine regressions across 77 countries, moving financial depth from the mean of the slowest-growing quartile (0.2) to that of the fastest-growing quartile (0.6) is associated with almost **1 percentage point of additional per-capita growth per year**. Against a 30-year gap of about five points between the slowest and fastest quartiles, financial depth alone accounts for a fifth (NBER, Levine, 2004). The instrumented estimates carry the same magnitude: had Argentina's private credit stood at the developing-country mean instead of 16 per cent of GDP, it would have grown **more than a percentage point faster per year** — in a period when it actually managed 1.8 per cent (Levine, Loayza & Beck, in NBER, Levine, 2004).

**Layer two: the causal channel.** Correlation could run backwards — rich economies buy deep finance. Rajan and Zingales closed that door: industries that are **structurally dependent on external finance** — pharmaceuticals, say, versus tobacco — grow disproportionately faster in financially developed economies. That differential pattern is what a causal finance-to-growth link

predicts and reverse causation cannot produce (NBER, Rajan & Zingales, 1996).

**Layer three: markets as well as banks.** Stock-market liquidity — the ease of trading claims — robustly predicts long-run growth, capital accumulation and productivity across countries, independent of banking depth (World Bank, Levine & Zervos, 1998). The steering system has two hands on the wheel.


**Layer four: the negative proof.** Hsieh and Klenow measured marginal products of capital and labour across manufacturing plants within narrow industries in China and India. In a well-steered economy those marginal products converge; instead they found dispersion of a factor of 6–7 between the 90th and 10th percentile plant, against under 2 in the US. Reallocating inputs to equalise marginal products would raise manufacturing TFP **on the order of a factor of two** — the published estimates for reaching US-level allocative efficiency are TFP gains of 30–50 per cent for China and 40–60 per cent for India (QJE, Hsieh & Klenow, 2009). Misallocation is a large fraction of why poor countries are poor. And the dial moves: China's own improvement over 1998–2004 — reallocation gains of roughly 2 per cent of TFP per year — coincided with exactly the market reforms Movement I credited.

| LAYER           | IDENTIFICATION                                               | HEADLINE RESULT                                                                  | SOURCE                           |
|-----------------|--------------------------------------------------------------|----------------------------------------------------------------------------------|----------------------------------|
| Cross-country   | 77-country growth regressions + instruments                  | Depth quartile move $\approx$ <b>+1pp</b> per-capita growth/yr                   | NBER, Levine 2004                |
| Causal channel  | External-finance-dependent industries $\times$ country depth | Dependent industries grow <b>disproportionately faster</b> where finance is deep | NBER, Rajan & Zingales 1996      |
| Markets channel | Stock-market liquidity across countries                      | Liquidity <b>predicts growth</b> independent of banks                            | World Bank, Levine & Zervos 1998 |
| Negative proof  | Plant-level marginal-product dispersion, China & India       | To US efficiency: TFP <b>+30–50% / +40–60%</b>                                   | QJE, Hsieh & Klenow 2009         |

## WHAT IT CONTRIBUTES – THE MECHANISM AT ITS SHARPEST

Venture capital is the specialised machinery for allocating capital to ideas with **no collateral, no cash flow and no history** — the projects banking cannot see. Kortum and Lerner, using three decades of industry panel data and the 1979 relaxation of the pension "prudent man" rule as a natural experiment, found a dollar of venture capital **three to ten times more effective** than a dollar of corporate R&D at producing patented innovation — so that VC, while less than 3 per cent of corporate R&D spending, plausibly accounted for about **15 per cent of US industrial innovations** in the decade to 1992 (NBER, Kortum & Lerner, 1998). The downstream footprint matches: of the 4,109 US IPOs between 1995 and 2018, **47 per cent were venture-backed**; VC-backed firms make up 56 per cent of the surviving cohort and account for **89 per cent of its R&D spending** — from a financing channel that touches under 0.5 per cent of new firms (HBS, Lerner & Nanda, 2020).

FIG. C4-1 – THE VENTURE FUNNEL: TINY CHANNEL, OUTSIZED OUTPUT



Less than half a per cent of firms carry nearly nine-tenths of their cohort's R&D — allocation quality, not capital quantity, is the scarce input. Sources: NBER, Kortum & Lerner 1998; HBS, Lerner & Nanda 2020.

## FAILURE MODE

Finance is the machine's most **accident-prone component**, with three documented faults. **Too much finance**: the IMF's evidence finds the growth benefit of deepening weakening — and stability risk rising — beyond a threshold, when finance shifts from allocating capital to trading claims on itself (IMF, Rethinking Financial Deepening, 2015). **Misallocation by politics**: the Hsieh–Klenow gaps are largely made of subsidised credit to connected and state-owned firms — finance captured by a failing institutions layer beneath it. **Concentration of the sharpest channel**: the venture model reaches a narrow slice of geography and founders — the top 25 US urban areas absorb about **75 per cent of all venture capital** (HBS, Lerner & Nanda, 2020) — leaving allocation blind spots the mechanism itself admits. All three are regulatory and structural design problems, in a component whose measured upside — a doubling of manufacturing productivity — is worth the engineering.

## WHY THE VALUE GOES TO SOCIETY

The financier's take is the spread or the equity gain; the social return is everything the capital built — the jobs, the products, rank 1's spillovers — none of which the intermediary can invoice. Venture fees are a microscopic fraction of the value of 15% of US industrial innovation, of which (per Nordhaus) the innovators themselves keep ~2%.

THE AGENTIC HOOK

Finance **is** information processing — the BIS describes the financial system as the brain of the economy, aggregating vast information into the price signals that coordinate everyone else (BIS, Aldasoro et al., 2024). Agentic underwriting extends credit to borrowers currently priced out by **assessment costs** and steers capital by cash-flow and marginal-product signals rather than collateral and relationships. That is the Hsieh–Klenow correction run through the credit channel — why Movement III counts allocation repair among the largest terms of the agentic dividend. The caveat, from the same BIS paper: correlated models can herd, so supervision must upgrade with underwriting.

PRINCIPLES

- 1

**Judge finance by allocation, not by size.** Beyond a measurable threshold, more finance buys instability, not growth — deepen the steering, not the balance sheet (IMF, 2015).
- 2

**Audit politically allocated credit annually.** The Hsieh–Klenow factor-of-two is what the audit prevents.
- 3

**Build the early-stage channel outside the capital city.** Three-quarters of US venture capital lands in 25 metros; the blind spots are where the unfunded challengers live (Lerner & Nanda, 2020).
- 4

**Open the data rails.** Underwriting quality is bounded by what lenders can legally see — e-invoicing, open banking and machine-readable registries are capital-allocation policy.
- 5

**Keep markets liquid as well as banks solvent.** The equity channel independently predicts growth (Levine & Zervos, 1998).

EXAMPLES

- ▶

**US pension rule change, 1979** — the "prudent man" clarification released pension money into venture funds: the natural experiment behind the 3–10× potency estimate (NBER, Kortum & Lerner 1998).
- ▶

**Argentina, counterfactual** — at developing-country-average private credit instead of 16% of GDP, growth >1pp/yr faster (NBER, Levine 2004).
- ▶

**US IPO cohort 1995–2018** — 47% venture-backed; 89% of the cohort's R&D from a <0.5%-of-firms channel (HBS, Lerner & Nanda 2020).

|                                                      |                                                                    |                                                        |                                                           |
|------------------------------------------------------|--------------------------------------------------------------------|--------------------------------------------------------|-----------------------------------------------------------|
| <b>+1pp</b>                                          | <b>15%</b>                                                         | <b>75%</b>                                             | <b>×2</b>                                                 |
| growth/yr per financial-depth quartile move (Levine) | of US industrial innovations from <3% of R&D spend (Kortum–Lerner) | of US VC concentrated in 25 metro areas (Lerner–Nanda) | manufacturing TFP forgone to misallocation (Hsieh–Klenow) |

## MOVEMENT II · THE MACHINE · MECHANISM DEEP-DIVE

## 5 Entrepreneurship — the entry-exit experiment

## IN SHORT

Every claim in ranks 1–4 requires someone to actually found the challenger firm. Entrepreneurship is the machine's **experimental protocol**: thousands of entrants embody guesses about technology and demand; the market runs the trial; most fail and release their resources; a few grow explosively and become the new incumbents. The value comes from the distribution, not the average entrant.

~3%

Of US employment sits in startups — the mechanism's footprint looks small...

NBER · Haltiwanger et al. 2010

~20%

...yet startups generate almost a fifth of all gross job creation, every year

NBER · Haltiwanger et al. 2010

40%

Of a startup cohort's jobs are gone within five years — the experiment's deliberate failure half

NBER · Haltiwanger et al. 2010

16→4%

Collapse of the high-growth skewness premium after 2000 — the warning light on this mechanism

NBER · Decker et al. 2015

## WHY IT RANKS HERE

Entrepreneurship ranks fifth not because entry matters little but because its value is **partly executed through the mechanisms above it**: the entrant's price cuts are rank 3's transfer, its funding is rank 4's allocation, its new product is rank 1's innovation. What entry contributes irreducibly is the **experiment itself** — the only known procedure for discovering which ideas deserve resources, run at the economy's expense of mostly-failures and occasionally spectacular successes.

The rank is earned by a finding that overturned the most durable folk claim in economic policy — that *small* business creates the jobs. Control for firm size and the size effect disappears entirely: **there is no systematic relationship between firm size and growth once age is held constant** (NBER, Haltiwanger, Jarmin & Miranda, 2010). What matters is youth. Policy aimed at smallness subsidises the average; the value lives in the young, high-variance tail.

“

*The value is generated by the distribution, not the average — which is why policy aimed at the average entrant misses the point.*

ENSI · THE MACHINE — REPORT 2 OF 3, 2026

## WHAT IT CONTRIBUTES

**The age finding, in full.** Haltiwanger, Jarmin and Miranda used the Census Bureau's Longitudinal Business Database — all US non-farm firms, 1976 to 2005 — to separate the effects of firm size and firm age for the first time at national scale. Startups hold only about **3 per cent of US employment but generate almost 20 per cent of gross job creation**; their entering cohort looms large against a net national job flow averaging 2.2 per cent a year. The dynamic is brutally selective — about **40 per cent of the jobs a startup cohort creates are gone within five years** through exit — but conditional on survival, young firms grow far faster than mature ones. This "up or out" pattern appears in every sector (NBER, Haltiwanger, Jarmin & Miranda, 2010).

**Replication across datasets and borders.** The Kauffman Foundation's indicator series confirms the age finding in recent US data (Kauffman, 2022). The OECD's DynEmp project, harmonising firm-level microdata across countries, confirms it internationally —

with a policy twist: the **national policy environment shapes how far young firms can climb**, which is why the same experiment yields different value in different jurisdictions (OECD, No Country for Young Firms, 2016). The Global Entrepreneurship Monitor tracks the base rate of entry across roughly 50 economies, giving the denominator: how many experiments each society runs (GEM, 2024).

**Why the value goes to society — three leaks.**

The **jobs**: net employment creation accrues to workers, not founders — and the founders of the failed 40 per cent keep nothing at all, while their experiments' findings (what does not work) inform every subsequent entrant free of charge. The **consumer surplus**: entrants win by offering better terms than incumbents — rank 3's transfer, executed by new firms. The **knowledge**: successful entry reveals demand nobody had measured — Hayek's discovery procedure, with company registration numbers.

FIG. C5-1 – SMALL FOOTPRINT, OUTSIZED JOB ENGINE



**Firms in their first years create jobs at nearly seven times their employment weight.** Age, not size, is the variable that matters. Source: NBER, Haltiwanger, Jarmin & Miranda 2010 (US LBD, 1976–2005).

**THE MYTH THIS BURIES**

"Small business creates the jobs" fails the age control. A forty-year-old ten-person firm is not a job engine; a two-year-old ten-person firm may be. Retire small-business policy; adopt **young-firm policy**.

## FAILURE MODE

The experiment is slowing down — and the character of the slowdown changed around 2000. Decker and co-authors document that US business dynamism has declined for decades; the earlier decline was broad, but the post-2000 decline is concentrated where it hurts most: the **high-growth tail is thinning**. In 1999, the firm at the 90th percentile of employment growth grew about 31 per cent faster than the median, and that 90–50 differential was 16 per cent larger than the 50–10 differential below it. **By 2007 the skewness premium had collapsed to 4 per cent**, and it kept falling through 2011 — driven both by a declining share of young firms in the economy and by a declining propensity of the young firms that do exist to be high-growth (NBER, Decker, Haltiwanger, Jarmin & Miranda, 2015). Fewer experiments, and tamer ones.

The diagnosis points upstream. The CEA links the long entry decline to rising concentration and to regulatory entry barriers such as occupational licensing (CEA, 2016) — meaning this fault is substantially **downstream of the rank-3 fault** and shares its toolkit: enforcement, entry-barrier demolition, licensing reform, plus insolvency speed so the "out" half of up-or-out recycles resources quickly.

FIG. C5-2 — THE THINNING TAIL



The excess growth of the fastest young firms over the median collapsed after 2000 — the plausible twin, seen from the other side, of the markup rise in Fig. C3-1. Source: NBER, Decker et al. 2015 (90–50 vs 50–10 employment-growth differentials).

## ONE FAULT, TWO GAUGES

The post-2000 skewness collapse (rank 5) and the markup rise (rank 3) are plausibly the **same fault observed twice**: challengers who cannot break through are also rents that never get competed away.

## SOURCE NOTE — TRANSPARENCY

Gornall & Strebulaev's estimates of VC-backed firms' share of public-market value could not be retrieved for this library (source gap; file removed as mis-downloaded). All VC-footprint figures in this report therefore cite **Lerner & Nanda (HBS, 2020)**: 47% of 1995–2018 IPOs, 89% of cohort R&D.

## THE AGENTIC HOOK

Why do would-be founders not found? Because a firm is a bundle — product, sales, accounting, compliance, support, logistics — and assembling the bundle has a minimum cost in people and coordination. Every function an agent performs at near-zero marginal cost lowers that threshold: the **minimum viable firm size is falling**. The field evidence reads as an entrepreneurship finding, not just a labour one — an agent carrying expert tacit knowledge lifted *novice* workers' productivity by 34 per cent (NBER, Brynjolfsson, Li & Raymond, 2023), which means an agent stack can hand a first-time founder the operational competence of a seasoned team. The one-person, agent-operated firm with meaningful revenue is a documented phenomenon in 2026. Two consequences: the **experiment rate rises** — the Decker decline run in reverse — and the composition of entry shifts toward the previously excluded: domain experts without operations skill, founders outside capital hubs and outside English. The state's lever is to make itself operable by the founder's agents: registration, tax and licensing as machine-consumable rails. **"Agent-operable government" is the new ease of doing business** — and it is measurable, the way B-READY measures regulatory efficiency today. Movement III specifies the build.

## PRINCIPLES

- 1 **Target age, not size.** Subsidies keyed to smallness reward staying small; the evidence keys value to youth and growth (Haltiwanger et al., 2010).
- 2 **Price the runway.** Startup costs, licensing burden and insolvency speed decide whether "up" is executable — the DynEmp finding is that policy sets young firms' ceiling (OECD, 2016).
- 3 **Protect the failure half of the experiment.** Fast, non-punitive exit recycles people, capital and the free knowledge of what does not work.
- 4 **Watch skewness, not entry counts.** A thousand lifestyle entrants matter less than one 90th-percentile grower; the 16→4 gauge is the one to publish (Decker et al., 2015).
- 5 **Make government agent-operable.** Company registration, tax and licensing exposed as machine-consumable services is entry policy for the agentic decade.

## EXAMPLES

- ▶ **US Longitudinal Business Database, 1976–2005** — the age-not-size result across all non-farm firms and every sector (NBER, Haltiwanger, Jarmin & Miranda 2010).
- ▶ **United States, post-2000** — high-growth skewness premium 16% → 4% by 2007; young-firm share and high-growth propensity both declining (NBER, Decker et al. 2015).
- ▶ **OECD DynEmp, cross-country** — same up-or-out physics everywhere; national policy environments set how far the "up" runs (OECD 2016).
- ▶ **~50 economies, annually** — GEM tracks the global base rate of entry: how many experiments each society runs (GEM 2024); Kauffman confirms the US age finding in current data (2022).

## MOVEMENT II • THE MACHINE • MECHANISM DEEP-DIVE

## 6 Trade — the machine, extended across borders

## IN SHORT

Trade is the five mechanisms above operating at planetary scale: prices coordinating across borders, competition arriving by ship, capital and ideas following the goods. Specialisation by comparative advantage raises the productivity of both parties without either inventing anything — the only mechanism in this catalog that creates value by **pure rearrangement**.

### \$2.1T/yr

US income payoff from trade expansion, 1950–2016 — one country's annual dividend

PIIE • Hufbauer & Lu 2017

### \$18,131

The same dividend per US household per year ( $\approx$  \$7,014 per capita)

PIIE • Hufbauer & Lu 2017

### 0.5–2%

Income gain per person for each +1 percentage point of trade share — geography-instrumented

Frankel & Romer • via NBER, ACR 2010

### ~1B

People lifted from extreme poverty since 1990, with trade openness credited as central

WTO & World Bank 2015

## WHY IT RANKS HERE

Trade ranks sixth **not because the dividend is small** — in absolute dollars it rivals anything in this catalog — but because it is **derivative**: it multiplies the value the other mechanisms create rather than creating value independently. A country that got ranks 0–5 right would still grow rich trading only moderately; no volume of trade rescues an economy whose domestic machine is broken — the resource-exporter counterexamples in Movement I are the proof by existence. The ranking criterion is substitutability, and trade is the most substitutable of the seven precisely because it is the same machine, wider.

The theory carries unusual authority for a policy debate this loud: Arkolakis, Costinot and Rodríguez-Clare showed that across an entire family of workhorse trade models, the **welfare gains from trade are pinned down by just two observable statistics** — the share of expenditure on domestic goods and the elasticity of trade flows (NBER, ACR, 2010). Different micro-stories, same sufficient statistics: the gains are robust to which model of trade you believe.

“

*No volume of trade rescues an economy whose domestic machine is broken — trade is the machine, wider.*

ENSI • THE MACHINE — REPORT 2 OF 3, 2026

## WHAT IT CONTRIBUTES

**The rich-country receipt.** The Peterson Institute's accounting puts the payoff to the United States from trade expansion between 1950 and 2016 at about **\$2.1 trillion a year** — roughly **\$18,131 per household** (PIIE, Hufbauer & Lu, 2017). Behind it stands the geography-instrumented causal estimate: Frankel and Romer's finding that a one-percentage-point rise in the trade share raises income per person by **0.5 to 2 per cent**, an estimate that remains standing in the modern quantitative-trade literature built on the ACR sufficient statistics (NBER, Arkolakis, Costinot & Rodríguez-Clare, 2010).

**The poor-country receipt.** The joint WTO–World Bank study credits trade openness as central to the defining welfare event of the era — the exit of roughly **a billion people from extreme poverty since 1990**, as poor-country producers connected to world demand (WTO & World Bank, 2015). The World Development Report on global value chains supplies the mechanism: GVC participation raises productivity, wages and formal employment in developing economies — the channel by which a Vietnamese or Bangladeshi worker joins the

world price for their labour rather than the local one (World Bank, WDR 2020).

**Why the value goes to society.** The gains from trade arrive overwhelmingly as **consumer surplus** — lower prices, wider variety — and as productivity forced on domestic firms by import competition (LSE CEP, Bloom, Draca & Van Reenen, 2011). Importers and exporters capture margins; the \$18,131 per household is the leak. Trade's leak rate is competition's, extended across borders.

**Failure mode: concentrated, slow-healing local losses.** The China Shock literature documents that import competition's costs fell on specific commuting zones, with wage and employment effects persisting a decade or more — about **2.0–2.4 million US jobs** over 1999–2011, real national surplus and real regional devastation at once (Autor, Dorn & Hanson, 2016). The engineering read from Movement I stands: the fault lies in the **compensation layer** — retraining, mobility, place-based policy, a fiscal institution of rank 0 — not in the trade mechanism itself. Tariff walls, the same literature shows, destroy far more surplus than they protect.

| EVIDENCE              | DESIGN                                        | HEADLINE RESULT                                                            | SOURCE                    |
|-----------------------|-----------------------------------------------|----------------------------------------------------------------------------|---------------------------|
| Sufficient statistics | Welfare formula across workhorse model family | Gains pinned by <b>two observables</b> : domestic share + trade elasticity | NBER, ACR 2010            |
| Geography instrument  | Predicted trade from distance/size            | <b>+0.5–2%</b> income per +1pp trade share                                 | Frankel & Romer           |
| National accounting   | US integration 1950–2016                      | <b>\$2.1T/yr</b> ≈ \$18,131/household                                      | PIIE, Hufbauer & Lu 2017  |
| Poverty channel       | Trade & poverty joint review; GVC microdata   | Trade central to the <b>~1B-person exit</b> ; GVCs raise wages & formality | WTO-WB 2015 • WDR 2020    |
| The debit             | Commuting-zone exposure to imports            | <b>2.0–2.4M</b> US jobs, decade-long local scarring                        | Autor, Dorn & Hanson 2016 |

## Kerala's fishermen: information becomes surplus

Between 1997 and 2001, mobile-phone service arrived in stages along the coast of Kerala, India. For the first time, fishermen at sea could call ahead and compare prices across beach markets before deciding where to land the catch. Jensen's micro-survey data record what followed — the price mechanism, repaired by one information technology, in one market, with every variable measured (QJE, Jensen, 2007).

### 5–8%→0

Share of the daily catch wasted — dumped unsold — before phones; eliminated after

QJE · Jensen 2007

### +8%

Fishermen's profits — the producer side gained...

QJE · Jensen 2007

### –4%

...while consumer prices fell at the same time. Both sides kept the difference

QJE · Jensen 2007

### ≈LOP

Price dispersion across markets collapsed; adherence to the Law of One Price became near-perfect

QJE · Jensen 2007

### WHAT THE CASE PROVES

Nobody was expropriated; the value came out of **dead-weight loss** — fish that once rotted found buyers, and the surplus split between producer and consumer (consumer surplus in sardine consumption rose 6 per cent). Information technology moved a real market measurably closer to the textbook ideal. That is the gains-from-exchange story of this whole mechanism, in one fishing coast, over four years.

### THE AGENTIC HOOK

The Kerala fisherman made **one phone call per trip**. A trading or consumer agent makes the equivalent of thousands per second, across every market it is pointed at. And trade's binding frictions are exactly agent-shaped: search, contracting, language, certification, customs — information and transaction costs, not tariffs. An export agent that finds demand in nineteen markets, drafts compliant contracts and files the customs declarations is a trade department rented by the hour — and it is **small-firm-biased**, because it commoditises the fixed costs only large firms could amortise (Microsoft Research, Rothschild et al., 2025). Movement III runs Kerala continuously, economy-wide.

### PRINCIPLES

- 1 Openness is cheap to keep and expensive to replace.** The \$2.1T dividend compounds silently; tariffs bill immediately (PIIE, 2017).
- 2 Fund the adjustment layer as trade policy.** The China Shock bill is the cost of having skipped it, not the cost of trade (Autor, Dorn & Hanson, 2016).
- 3 Cut fixed costs, not just tariffs.** The binding frictions are informational — customs, certification and contracting are where the next dividend sits.
- 4 Measure by household, not by trade balance.** \$18,131 per household per year is the statistic that describes who trade is for.

### EXAMPLES

- ▶ **United States, 1950–2016** — integration worth \$2.1T/yr, ≈\$18,131/household (PIIE 2017).
- ▶ **Kerala, India, 1997–2001** — phones erase price dispersion; waste to ~zero; profits +8%, prices –4% (Jensen 2007).
- ▶ **GVC entrants — Vietnam, Bangladesh** — value-chain participation raises productivity, wages, formality (WDR 2020).
- ▶ **The billion-person exit** — trade openness credited as central since 1990 (WTO-WB 2015).

### BRIDGE

Kerala is what happens when **one** information technology repairs **one** mechanism **once**. Movement III asks what happens when agentic AI repairs all seven, continuously.

## MOVEMENT II • THE MACHINE • SYNTHESIS

# Σ The compounding stack

Ranking the mechanisms separately, as this movement just has, **understates them jointly** — because the strongest results in the library are about the couplings.

## Competition × Innovation

Rivalry is a technology-forcing device: Chinese import competition raised European firms' patenting, IT adoption and productivity (LSE CEP, Bloom, Draca & Van Reenen 2011), and weak competition is a principal cause of the badly-managed tail (Bloom & Van Reenen 2006). **Rank 3 feeds rank 1.**

## Finance × Entrepreneurship

Up-or-out needs someone to fund the "up": venture capital exists to finance rank 5's high-variance young firms, and its 3–10× per-dollar potency (Kortum & Lerner 1998) is earned on exactly the firms that create the disproportionate jobs and R&D (Lerner & Nanda 2020). **Rank 4 feeds rank 5.**

## Prices × Everything

The information layer lets every other mechanism act at a distance: the financier reads risk in spreads, the entrepreneur reads unmet demand in margins, the trading firm reads comparative advantage in relative prices. Hayek's tin cascade (AER, 1945) is the protocol they all run on. **Rank 2 carries the signals.**

## Institutions × Everything

The Hsieh–Klenow decomposition is the negative demonstration: China's and India's factor-of-two shortfall is *finance* misallocating because *institutions* override *prices* — three mechanisms failing as one coupled system (QJE, 2009). China's reform decades show the couplings running forward — reallocation gains ≈ 2% of TFP/yr. **Rank 0 hosts the stack.**

### ONE FAULT SEEN FROM TWO SIDES

Entry is how creative destruction actually executes: the challenger that obsoletes the incumbent is, empirically, a young firm. The post-2000 thinning of the high-growth tail (Decker et al. 2015) and the 18→67% markup rise (De Loecker & Eeckhout 2017) are plausibly the **same degradation observed on two gauges.**

The compounding is the deep reason referendum-style politics mishandles capitalism: **you cannot keep the innovation dividend while suspending competition**, or keep entrepreneurship while letting finance be politically allocated. The correct policy posture is the maintenance of a machine, not the adjudication of an ideology.

“

*The machine's components are load-bearing for each other.*


ENSI • THE MACHINE — REPORT 2 OF 3, 2026

## MOVEMENT II • THE MACHINE • THE ASSEMBLY SEQUENCE

## How the stack assembles — and why the whole exceeds the parts

FIG. C7-1 — THE VALUE ENGINE, ASSEMBLED BOTTOM-UP

ASSEMBLY ORDER: INSTITUTIONS → PRICES → COMPETITION → FINANCE → ENTRY → INNOVATION → TRADE



**The build order is the argument:** each layer only works when the layers below it stand — which is why the counterfactual gradients (AJR ~7×, Fraser 7.6×) dwarf any single mechanism's annual dividend. Sources: MIT AJR 2001; AER Hayek 1945; NBER Jones & Summers 2020; assembly reading — ENSI.

The whole exceeds the parts because each mechanism's measured dividend was estimated **with the others running**. Jones and Summers' \$13.30 assumes competition passes gains to users and finance funds the labs; Levine's +1pp assumes prices carry honest signals; Haltiwanger's job engine assumes entrants can borrow and incumbents can be displaced. The stack's arithmetic is multiplicative, and the leak — 98 per cent of the surplus escaping to society — is the emergent property of all seven layers operating at once.

## SEGUE → MOVEMENT III

Every layer of this stack is, at bottom, a machine for processing information and lowering the cost of transacting. **Agentic AI is a general-purpose collapse in the price of exactly those two things** — the first component in a century that upgrades every other part of the stack at once. Movement III installs it.

## MOVEMENT III – THE MULTIPLIER

# Agents don't run on the machine. They upgrade it.

Every mechanism ranked in Movement II is, at bottom, a machine for processing information or lowering the cost of transacting. Agentic AI is a general-purpose collapse in the cost of exactly those two things — so it does not add output beside the seven mechanisms; it multiplies through them. This movement maps the amplification mechanism by mechanism, prices the conditional dividend transparently, and states the guardrail conditions under which society — not the model layer — keeps it.

III.1 The Amplification Thesis & the Evidence Floor – p. 96

III.2 Amplification Map I: Prices · Competition · Innovation – p. 99

III.3 Amplification Map II: Entrepreneurship · Capital · Trade · Institutions – p. 103

III.4 The Agent-Native Economy – p. 106

III.5 The Nordhaus Question for AI – p. 109

## MOVEMENT III • THE MULTIPLIER

## III.1 The Amplification Thesis

Look again at the seven mechanisms of Movement II through one lens: every one of them is an information-processing or transaction-cost machine. That is not a metaphor — it is Hayek's and Coase's original arguments, and it is why agentic AI is not one more productivity tool.

Hayek's point about the price system was never that markets are virtuous but that they are **the only computer big enough**: a distributed computation over knowledge no participant possesses in full. Competition works only where consumers can actually compare offers — search costs and obfuscation are its friction. Innovation is the production of ideas, whose main input is cognitive labour. Entrepreneurship is bounded by how much coordination one founder can afford — Coase's transaction costs are what makes firms necessary and small teams weak. The BIS calls capital allocation the brain of the economy: "the processing and aggregation of vast amounts of information into price signals that coordinate participants in the economy." Trade is limited less by tariffs than by the costs of search, contracting, language and coordination across borders. And institutions — North's humanly devised constraints — are priced in the apparatus of verifying, enforcing and complying with rules.

### Information machines

The **price system** aggregates dispersed knowledge; **innovation** produces ideas from cognitive labour; **capital allocation** converts information into credit and prices. Agents are elastic cognitive labour: reading, comparing, computing without fatigue.

### Transaction-cost machines

**Competition** is throttled by search costs; **entrepreneurship** by the cost of assembling a firm; **trade** by export fixed costs; **institutions** by the cost of verification and compliance. Agents transact, negotiate and verify at near-zero marginal cost.

If that reading is right, a technology that collapses the cost of information-processing and transacting does not *add to* the seven mechanisms the way a new industry adds output. It is an input into the machinery that produces all the other gains — it **multiplies through** every mechanism at once. That is the amplification thesis, and it is why the Goldman–Acemoglu gap of Analysis 11 (+7% of global GDP versus  $\leq 0.71\%$  of TFP) is not a disagreement about coefficients: it is a bet on whether AI stays a tool applied to fixed tasks or becomes an actor inside the mechanisms themselves.


### THE DISCIPLINE RULE FOR THIS MOVEMENT

Every amplification claim below must clear two bars: a **measured floor** from field evidence, and a **stated condition** under which the gain fails or is captured. An amplification thesis that cannot state its own failure conditions is marketing.

## III.1 • THE EVIDENCE FLOOR

## What the field experiments already show

The base of the argument is no longer speculative: real deployments already show double-digit gains — and every number below was produced by the **weakest possible form of agency**: a suggestion tool.



80%

of US workers have ≥10% of tasks exposed to LLMs (Eloundou et al. 2023)

19%

could see at least half of their tasks affected (same rubric)

15%→47–56%

share of all tasks completable significantly faster: model alone → with agent tooling built on top

The anchor study is a staggered rollout of a generative-AI assistant to **5,179 customer-support agents**. The average +14 per cent conceals the finding that matters: novices gained 34 per cent while the most experienced gained almost nothing, because the model captures the tacit knowledge of the best performers and disseminates it. Customer sentiment improved; retention rose. AI assistance **compressed skill inequality inside the firm**.

The laboratory evidence points the same way on different tasks: writing time down 37 per cent while grades rose 0.45 standard deviations, with low-graded workers improving on both margins — and work restructuring away from rough-drafting toward idea-generation and editing.

The quiet hinge of the macro debate sits in the exposure rubric Acemoglu himself builds on: with the model alone, about 15 per cent of worker tasks could be done significantly faster at equal quality; **with software and tooling built on top, 47–56 per cent**. That leap — from a sixth of tasks to half — is precisely the leap from model to agent. Goldman's occupational analysis lands in the same territory: roughly two-thirds of jobs exposed, some 300 million full-time-equivalent roles globally.


The discipline note: **none of these results yet involves agents transacting**. They are a floor, established with the weakest agency deployed. Everything the rest of this movement adds sits on top of — and must be argued from — this measured base.

## III.1 • THE EVIDENCE FLOOR

## Usage is fact, and it is now measured

Exposure is potential. The Anthropic Economic Index turns usage into telemetry: over four million Claude.ai conversations mapped onto the US Department of Labor's O\*NET task taxonomy.

FIG. D-1 • HOW AI IS USED, BY INTERACTION PATTERN



The 43 per cent automation share is the leading indicator of the tool→actor transition — the number to re-read quarter by quarter.

~36%

of occupations already use AI on at least a quarter of their associated tasks

Anthropic Economic Index • Handa, Tamkin et al. 2025

~4%

of occupations use AI on three-quarters or more of tasks — deep automation of whole jobs is still rare

Anthropic Economic Index 2025

The telemetry shows an economy in the transition zone. Software development and writing account for nearly half of all usage, but diffusion is already broad: a third of occupations use AI on a quarter of their tasks within three years of the technology existing. Usage peaks in the upper-middle of the wage distribution and falls at both extremes. AI is performing economic tasks at scale — mostly alongside humans, increasingly instead of them.

Read as a dashboard rather than a debate, the evidence floor gives the amplification thesis its empirical tripwires. If the automation share stalls, if the profitably-automatable fraction of exposed tasks stays near Svanberg's 23 per cent, and if measured task portfolios never contain tasks absent from O\*NET, then Acemoglu's bound holds and this movement overclaims. If the automation share climbs and new task categories appear, the actor scenario is arriving through measurable channels.

#### FLOOR, NOT CEILING — AND NOT DESTINY

Every number on this spread was generated by assistants, not transacting agents. That cuts both ways: the upside case is larger than the measurements, and **so is the unmeasured risk surface**. The rest of this movement carries both.

Anthropic EI 4M+ conversations • 2025

O\*NET task taxonomy

Acemoglu MIT 2024 • tripwires

## MOVEMENT III • THE MULTIPLIER

## III.2 Amplification Map I: Prices, Competition, Innovation

### Prices — the continuous Kerala

−5–8%

Daily catch previously wasted — eliminated entirely once fishermen could compare beach prices by phone

QJE • Jensen 2007

−4%

Consumer price of fish — while adherence to the Law of One Price became near-perfect

QJE • Jensen 2007

+8%

Fishermen's profits — both sides of the market kept the difference; the value came out of dead-weight loss

QJE • Jensen 2007

The price system's failure mode is specific: where information is costly, arbitrage stops, dispersion persists, and goods rot in the wrong markets. That failure mode has been dissolved by an information technology once before, and it is the cleanest natural experiment in this library. Between 1997 and 2001, mobile-phone coverage arrived along the Kerala coast; fishermen at sea could call ahead and compare market prices before landing the catch. Dispersion collapsed, waste went to zero, profits rose **while** consumer prices fell. Nobody was expropriated. **Information technology moved a real market measurably closer to the textbook ideal, and both sides kept the difference.**

The Kerala fisherman made one phone call per trip. A consumer or firm agent makes the equivalent of thousands per second, across every market it is pointed at, without fatigue, pride or satisficing. The agentic amplification of the price mechanism is Jensen's result **run**

**economy-wide and continuously:** agents monitoring, comparing and arbitraging prices for electricity, freight, insurance, credit, software and professional services — exactly the categories where dispersion is wide today because human search is costly and episodic.

The BIS frames finance's information capacity as "shaped in large part by the information processing technology available" — Sumerian book-keeping, double-entry accounting, the telegraph, AI. Agents are the next instalment in that series, with one difference of kind: every previous technology made human arbitrageurs faster. **Agents are the arbitrageurs.**

The policy conversion is precise. Once agents are allowed to see and act on prices, dispersion should collapse as it did in Kerala. Where it persists — healthcare, legal services, public procurement — **the residual dispersion is a live map of protected rents**, updated continuously, sector by sector.

#### THE FORESIGHT INDICATOR

Not "AI adoption" — **the dispersion statistics themselves**. A state that publishes sectoral price-dispersion indices gets a rent map for free, and a running test of whether the amplification is real or blocked.

## III.2 • AMPLIFICATION MAP I

**Competition — the end of the search-cost moat**

A large share of modern market power lives not in patents or scale but in **friction**: consumers who cannot compare, switching costs deliberately inflated, drip pricing, obfuscated fee schedules, loyalty exploited because re-searching the market costs an evening. The counterfactual of unpoliced friction is already on the books — average US markups rose from 18 per cent above marginal cost in 1980 to 67 per cent in 2014 (De Loecker & Eeckhout).

Microsoft Research's analysis of the agentic economy identifies exactly this as the deep effect: the profound impact of generative AI lies in reducing "communication frictions between consumers and businesses", not in speeding up existing workflows. Their canonical example is mundane and devastating: a consumer stays with a worse tax preparer because explaining her financial situation again is too costly. An assistant agent that carries her context, interrogates every rival's service agent and re-runs the comparison monthly cuts that cost to roughly zero. **Obfuscation pricing does not survive a counterparty that reads every footnote; search-cost moats do not survive a shopper with infinite patience.** Every dark pattern in the e-commerce playbook is a bet on human attention limits, and agents void the bet.

But the same paper carries the conditional on which this whole movement turns. **Unscripted** interaction (agents *can* talk to anything) is not **unrestricted** interaction (agents are *allowed* to transact across firm boundaries). The architecture fork:

**"Agentic walled gardens"**

Inter-agent commerce runs through a few dominant providers, like today's app stores. The agent that shops for you is owned by the platform you shop on. The search-cost moat is not destroyed — it is **nationalised by the gatekeeper**, and the markup curve bends further up.

**The open "web of agents"**

Agents freely connect and transact across boundaries, like the web itself. Competition intensifies toward the Bertrand textbook; obfuscation rents, loyalty rents and search rents are competed away at machine speed — to consumers.

“

*The competition mechanism is amplified in both worlds; the question is for whom.*

ENSI SYNTHESIS OF MICROSOFT RESEARCH • ROTHSCHILD ET AL. 2025

Who captures the surplus is therefore a **policy outcome, not a technical one**. Interoperability mandates, agent-access rights to published offers, data-portability instruments extended to agent-readable form: protocol policy is the competition policy of the 2030s.

## III.2 • AMPLIFICATION MAP I

Innovation — attacking the 18<sup>x</sup>-researchers problem**18<sup>x</sup>**

more researchers needed today than in the early 1970s to sustain Moore's-Law doubling (Bloom, Jones, Van Reenen & Webb, AER 2020)

**\$5–13.30**

social benefit per \$1 of R&D — conservative to baseline (Jones & Summers, NBER 2020)

**67%**

social internal rate of return on R&D vs 10–15% private — the engine is under-fuelled

Innovation carries the largest coefficients in this library — and the most troubling stylised fact in growth economics. Research productivity is declining sharply everywhere Bloom and co-authors look: aggregate growth is sustained only by throwing exponentially more researchers at the problem. In their accounting, growth equals research productivity times the number of researchers. The first term is falling; the second cannot rise forever, because **human researchers are the binding input**.

Agentic AI attacks both terms at once, and this is where the multiplier genuinely compounds rather than adds. Agents that read the entire literature, generate and triage hypotheses, run simulations, write and debug code and draft papers are, functionally, **additional researchers whose supply is elastic in compute rather than in demographics**. And to the extent they raise the productivity of human researchers — the Noy–Zhang result applied to

the knowledge-production function itself — they lift the falling term too.

This is also the channel Acemoglu's 0.55–0.71 per cent bound **explicitly excludes**: he sets aside AI's "revolutionary effects" on the process of science as unlikely to bind within ten years. That exclusion was defensible in April 2024. By July 2026 — with agentic research pipelines operating in materials discovery, protein engineering and mathematics — it is the single assumption in his paper most worth stress-testing, because it is the one with the largest multiplier behind it.

Why the multiplier compounds: if \$1 of research effort yields \$5+ of social benefit, an agent layer that cuts the effective cost of a unit of research does not save research budgets — it applies the 5:1 ratio to a cheaper input, multiplying the social return on every research dollar, public and private. The next page prices that arithmetic — and labels every step.

**THE INVERSION OF R&D POLICY**

The classic instrument — subsidise research because private actors under-invest in a 67% social IRR — remains. The new instrument is to subsidise the **agentic research infrastructure itself**: open scientific data in agent-usable form, compute for public research agents, evaluation testbeds. A dollar there raises the productivity of every subsequent research dollar.

III.2 • AMPLIFICATION MAP I

The research dividend, priced — and labelled

The report’s rule: back-of-envelope arithmetic is shown in full, every input sourced, every estimate labelled. This is the largest single term in the agentic dividend, so it gets the fullest audit.

| STEP | INPUT                                                                                                               | VALUE            | SOURCE / STATUS                                                                                     |
|------|---------------------------------------------------------------------------------------------------------------------|------------------|-----------------------------------------------------------------------------------------------------|
| 1    | Global R&D spend, 2025                                                                                              | ≈<br>\$2.9T/yr   | Library aggregate [04]                                                                              |
| 2    | Agentic cut in the effective cost of a unit of research (literature, hypothesis triage, simulation, code, drafting) | 20–30%           | ENSI estimate (BOE) — anchored to the –37% task-time floor, discounted for non-automatable lab work |
| 3    | Social benefit per \$1 of R&D                                                                                       | \$5 –<br>\$13.30 | Jones & Summers, NBER w27863 (conservative / baseline)                                              |
| 4    | Effective-R&D-base expansion × social multiplier                                                                    | +\$3–<br>8T/yr   | ENSI estimate (BOE) — conditional, decade horizon                                                   |

FIG. D-2 • THE RESEARCH CHANNEL, BOUNDED



Even the conservative bound exceeds most countries’ GDP. Both bars are labelled ENSI estimates: arithmetic on published coefficients, not measurements.

The honest statement, in full: **nobody has yet measured agentic R&D’s effect on research productivity at the economy level.** The paper that does will be the most important growth-economics result of the decade. Until it exists, this channel is a sized hypothesis with strong priors — the field-experiment floor on cognitive work, the 5:1 social-return ratio, and visible agentic pipelines in materials, proteins and mathematics — not a fact.

WHAT WOULD FALSIFY IT

If agent-assisted research groups show no measurable lift in output per researcher by ~2030; if the cost cut lands only on writing-up rather than discovery; or if bottlenecks simply move to experiments and equipment — then step 2 fails and this term shrinks toward zero. Track it; don’t assume it.

## MOVEMENT III • THE MULTIPLIER

## III.3 Amplification Map II: Entrepreneurship, Capital, Trade, Institutions

### Entrepreneurship — the one-person agentic firm

~3% → ~20%

startups' share of employment vs share of gross job creation (Haltiwanger, Jarmin & Miranda)

16% → 4%

collapse of the high-growth "skewness premium" among young firms post-2000 (Decker et al.) — the renewal rate slowing

+34%

novice productivity lift from an agent carrying expert tacit knowledge (NBER w31161)

Why do would-be founders not found? Because a firm is a bundle of functions — product, sales, accounting, compliance, support, logistics — and assembling the bundle has a minimum cost in people and coordination. This is Coase's transaction-cost boundary in practice: below a certain team size the founder does everything badly; hiring past it consumes the capital that should buy experiments. **Every function an agent performs at near-zero marginal cost lowers that threshold** — customer support, bookkeeping, contract drafting, localisation, ad operations, tax filing are becoming a subscription. The operational entry cost of a firm is heading toward approximately zero; the one-person, agent-operated firm with meaningful revenue is already documented in 2026, and the ten-person firm that behaves like a hundred-person firm is the economically significant version.

Read the Brynjolfsson–Li–Raymond result as an entrepreneurship finding rather than a labour-

market one: an agent stack gives a first-time founder **the operational competence of a seasoned team**. Two consequences follow. First, the experiment rate rises — entrepreneurship is a search over mostly-failures, so society's option value scales with attempts. This is the 16%→4% skewness decline run in reverse: the warning gauge of Movement II becomes the recovery gauge of the agentic decade. Second, the composition of entry shifts toward the previously excluded — domain experts without operations skill, founders outside capital hubs, founders in small language markets.

The condition, as always: the stack must exist beyond English, and the state's own interfaces must be usable by machines. **"Agent-operable government" is the new ease-of-doing-business** — and it is measurable the way the World Bank's B-READY framework measures regulatory efficiency today.

NBER Haltiwanger et al. • age not size

NBER Decker et al. 2015

World Bank B-READY 2024

## III.3 · AMPLIFICATION MAP II

## Capital — shrinking the misallocation tax

**+30–50%**

Manufacturing TFP gain in China from reallocating capital and labour to US-observed efficiency

QJE · Hsieh & Klenow 2009

**+40–60%**

The same reallocation gain in India — misallocation is one of the largest sources of forgone income on the planet

QJE · Hsieh & Klenow 2009

“

*At the core of the financial system is the processing and aggregation of vast amounts of information into price signals that coordinate participants in the economy.*

BIS WORKING PAPER 1194 · ALDASORO, GAMBACORTA, KORINEK, SHREETI & STEIN, 2024

Finance *is* information processing — the BIS’s framing, not ours — and every step-change in information technology, from double-entry book-keeping onward, has restructured it. The paper works through the four functions, and in each the agentic gain has the same shape. **Intermediation:** credit decisions on richer data at lower cost extend lending to borrowers currently priced out by assessment costs — the small firm whose loan appraisal costs more than the margin it would generate. **Insurance:** underwriting and claims collapse from weeks to minutes. **Asset management:** private-bank quality at retail cost. **Payments:** the machine-native rails of page 107.

Agent underwriting is the Hsieh–Klenow correction run through the credit channel: capital flowing to marginal-product signals rather than to collateral and relationships. Even in advanced economies, the gap between well-managed and badly-financed firms is the same tax at smaller scale. The state’s lever is not to run credit agents; it is **the data architecture that determines what agents can see** — open banking, e-invoicing, machine-consumable registries. A compact economy with mandatory e-invoicing could make itself the cheapest place in Europe to underwrite a small firm, a comparative advantage worth more than most subsidy programmes.

## THE BIS’S OWN FAILURE MODES — CARRIED, NOT WAVED PAST

Herding when many institutions run correlated models; opacity; and dependency on a few AI providers as a **new category of critical infrastructure**. The same paper that supplies the upside supplies the supervision agenda — flash-crash dynamics generalised beyond finance.

NBER Levine · depth quartile ≈ +1pp growth

NBER Kortum & Lerner · VC 3–10× patent potency

BIS WP 1194 · 2024

## III.3 • AMPLIFICATION MAP II

**Trade — the export department, rented by the hour**

The binding constraint on who trades has never been tariffs alone; it is the **fixed costs of exporting** — finding counterparties, language, contract law across jurisdictions, customs paperwork, certification, payment risk. Those fixed costs are why exporting is dominated by large firms with compliance departments. Every one of them is an information or transaction cost, which is to say: agent-soluble. An export agent that identifies demand in nineteen markets, drafts compliant contracts in the counterparty's language, files the customs declarations and monitors payment is a trade department rented by the hour. The distributional consequence matters most for small open economies: **agentic trade intermediation is small-firm-biased**, because it commoditises exactly the fixed costs only large firms could previously amortise. The prize is priced by the Frankel–Romer coefficient — every percentage point of trade share bought by lower fixed costs returns a multiple in income; the US baseline from trade expansion runs \$2.1T a year, about \$18,131 per household.

**Institutions — ambient compliance, and the gap**

Institutions create value by making promises enforceable, and their price is the apparatus of verification, enforcement and compliance — text, procedure and checking: the home turf of language-model agents. The amplification runs in three layers. **Compliance**: the fixed cost of knowing and following the rules, which falls disproportionately on small firms, becomes an agent service — regulation consumed by machines rather than interpreted by lawyers. **Enforcement**: agents that monitor performance and flag breaches early shrink the lag that Haggard and Tiede identify as the growth-relevant core of the rule of law. **Reach**: societies enforce only the rules they can afford to enforce; when verification costs collapse, the enforceable set expands — **better institutions become cheaper to have**.

**THE GOVERNANCE-GAP ASYMMETRY**

**Agents consume institutions faster than institutions adapt.** Who is liable when an autonomous agent transacts badly? What is the legal identity of a negotiating agent? DeepMind and Microsoft converge on the same missing layer — verifiable credentials, agent identity, audit trails — and no jurisdiction had shipped a complete answer as of July 2026. The gap is not an argument against the multiplier; it is the multiplier's dependency.

The upside of the same asymmetry: the first legal systems to make agent commerce *safe* — identifiable, insurable, adjudicable — capture the institutional franchise, as Delaware did for the corporation and English law did for global contracting.

PIIE Hufbauer &amp; Lu 2017

NBER Frankel &amp; Romer 1999

JEP North 1991

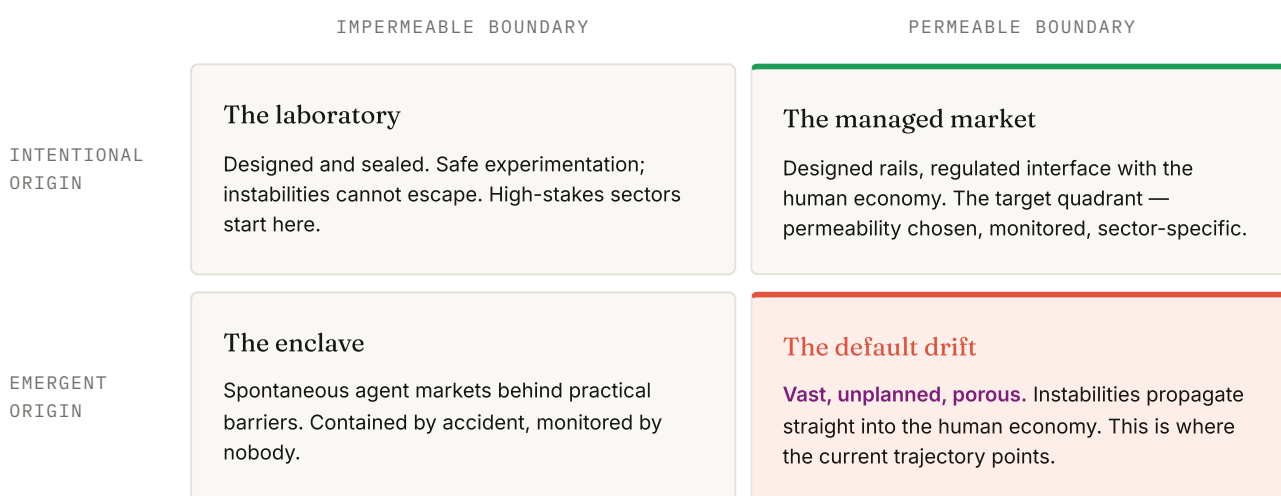
Haggard &amp; Tiede 2011

## MOVEMENT III • THE MULTIPLIER

## III.4 The Agent-Native Economy

Everything so far amplifies mechanisms that already exist. The final layer is new in kind: agents transacting *with each other*, at machine speed, in volumes and granularities no human economy could sustain — what Google DeepMind names the **sandbox economy**, a layer where agents transact and coordinate “at scales and speeds beyond direct human oversight.”

FIG. D-3 • DEEPMIND’S SANDBOX-ECONOMY QUADRANTS — AND WHERE THE DEFAULT LANDS



Axes from Google DeepMind, *Virtual Agent Economies* (Tomašev et al., 2025): origin (intentional vs spontaneous emergence) × boundary permeability. **The dangerous quadrant is the one nobody chooses — it happens by default.**

DeepMind’s central warning is a coordination-failure diagnosis, and it is worth quoting: “our current trajectory points toward a spontaneous emergence of a vast and highly permeable AI agent economy.” Permeability, the paper stresses, is **a collective property no single actor controls** — lowering it requires solving a collective-action problem, which is why drift, not design, is the default. Their agenda follows: auction mechanisms for fair resource allocation among agents, “mission economies” that point agent swarms at collective goals, and the trust infrastructure of verifiable credentials.

Google DeepMind Tomašev et al. 2025 • verified verbatim

## III.4 • THE AGENT-NATIVE ECONOMY

**Machine-native money, and value beyond human bandwidth**

The plumbing for the agent layer is being built now. The x402 protocol — initiated by Coinbase — resurrects the long-reserved HTTP 402 “Payment Required” status code as an open standard that “enables AI agents and web services to autonomously pay for API access, data, and digital services”: stablecoin settlement in ~200 milliseconds, no API keys, no subscriptions, no human checkout. The economics hiding in that engineering are large. Legacy rails carry per-transaction costs and chargeback risk that make anything below roughly a dollar uneconomic — which is why the internet’s business models collapsed into advertising and subscriptions. Machine-native micropayments reopen the design space Microsoft’s authors call unbundling and rebundling: paying per article, per API call, per second of compute, per answer. **A million transactions of a tenth of a cent is not a smaller version of a thousand transactions of a dollar; it is a different economy**, with different market structures, operating below the threshold of human attention.

What lives in that economy? Transactions that were always worth making and never worth a human’s time:

- ▶ **Procurement:** an agent collects and compares **10,000 supplier quotes before lunch** — a task no purchasing department ever priced, because no purchasing department could do it.
- ▶ **Data:** per-query licensing of a dataset — a market that cannot exist at subscription granularity, priced per question at machine speed.
- ▶ **Insurance:** per-hour cover on a single shipment, underwritten and settled in minutes — the BIS’s weeks-to-minutes collapse applied to a contract too small to exist today.
- ▶ **The dormant surplus:** every unswitched tariff, unclaimed refund, unexploited price difference and unread contract clause — executed instead of forgone.

A human consumer executes perhaps a few thousand economic decisions a year; almost everything else is surplus left on the table because **attention is the scarcest input in the consumer economy**. Agents convert that dormant decision-space into executed transactions. The aggregate is not a productivity statistic; it is the gap between the demand curve as economists draw it and the demand curve as exhausted humans actually walk it. None of it appears in a Hulten-theorem calculation — and none of it is governed yet, which is why the quadrant chart opposite comes first.

**x402 whitepaper** Reppel et al. 2025 • verified verbatim

**Microsoft Research** unbundling • 2025

**BIS** WP 1194


III.4 • THE AGENT-NATIVE ECONOMY

The conditional dividend: \$10–20 trillion a year, assembled in the open

The headline number of this movement is built from four channels, each shown with its arithmetic and its label. It is a **sized conditional, not a forecast** — and it deliberately excludes the agent-to-agent layer, which is real but unpriceable today.

| CHANNEL                       | MECHANISM AMPLIFIED      | ARITHMETIC                                                                                                        | ESTIMATE /YR                |
|-------------------------------|--------------------------|-------------------------------------------------------------------------------------------------------------------|-----------------------------|
| Research agents               | Innovation               | 20–30% effective research-cost cut × Jones–Summers \$5–13.30 social return per \$1 (p. 102)                       | +\$3–8T · ENSI estimate     |
| Price compression             | Prices · competition     | Kerala-scale dispersion collapse (prices –4%, waste –5–8%) run across ~\$60T global consumption                   | +\$1–2.4T · ENSI estimate   |
| Misallocation repair          | Capital allocation       | Agent underwriting closes half the Hsieh–Klenow gap on ~\$40T of emerging-market output — double-digit TFP points | +\$4T+ · ENSI estimate      |
| Entry & trade                 | Entrepreneurship · trade | Skewness premium re-widening from 4% toward 16%; commoditised export fixed costs (Frankel–Romer coefficient)      | positive, unquantified      |
| SUM, CENTRAL, WITHIN A DECADE |                          |                                                                                                                   | \$10–20T/yr · ENSI estimate |

FIG. D-4 • THE DIVIDEND BY CHANNEL (CENTRAL VALUES, \$T/YR)



Every bar is a labelled ENSI estimate (BOE): arithmetic run on published, peer-reviewed coefficients — Jones–Summers, Jensen, Hsieh–Klenow — not on model-vendor projections.

WHAT THE NUMBER IS – AND IS NOT

It is what the published coefficients imply **if agents become actors and competition survives**. It is not a prediction: the tool scenario (Acemoglu’s ≤0.71% TFP) and the captured scenario (multiplier fires, split flips — p. 109) both remain live. The conditions are the subject of the next section.

## MOVEMENT III • THE MULTIPLIER

## III.5 The Nordhaus Question for AI

Will agentic producers keep 2 per cent of the surplus — or 40? The whole report turns on Nordhaus's 2.2 per cent, so the honest question is whether the agentic century repeats it. The answer starts with an uncomfortable fact: **the 2.2 per cent was never a law of nature.**

Producers kept a minuscule fraction of innovation's value in the twentieth century because **competition and imitation relentlessly passed the surplus to consumers** — and competition and imitation were themselves maintained: by antitrust enforcement, by patent expiry, by open science, by entry barriers kept low. The 2.2 per cent is a **policy achievement**, the compound interest on a century of competition maintenance. Where the maintenance lapsed, the split drifted — and it drifted before a single agent transacted: average US markups rose from 18 per cent above marginal cost in 1980 to 67 per cent in 2014. The pre-agentic economy already shows what four decades of unpoliced market power do to the giveaway.

Now place the agent layer on top. Frontier models are supplied by a handful of firms; the walled-garden architecture is the default commercial instinct of every platform; and

Acemoglu adds that even benign AI is predicted to widen the gap between capital and labour income. If the model layer concentrates, **the multiplier still fires — but the split flips**. A concentrated agent layer levies a private tax on the price mechanism, on competition, on research and on trade simultaneously, because the same chokepoint sits inside all of them — a chokepoint Nordhaus's innovators never had. Poured into that structure, the most powerful surplus-creation technology ever built becomes **the most efficient rent collector ever deployed**.

Both futures run on identical technology. The Kerala-at-scale economy and the rent-collector economy differ in exactly one variable: whether the agent layer looks like the open web or like the app store. That is a choice, it is being made now, and it is being made mostly by default.

FIG. D-5 • THE DRIFT THAT PRECEDED THE AGENTS — AVERAGE US MARKUP OVER MARGINAL COST



The warning light was already on before the agentic era began. De Loecker & Eeckhout, NBER 2017. Coral marks the adverse trend, not an ENSI judgement.

## III.5 • THE NORDHAUS QUESTION FOR AI

**The guardrail conditions — under which the multiplier keeps its sign**

None of these is exotic. Each is ordinary policy of the kind that produced the original 2.2 per cent — applied to the agent layer while its equilibrium is still unformed.

- 1 **Open-agent access.** Any offer published to humans must be lawfully accessible to their agents on non-discriminatory terms. One clause, aimed squarely at the walled-garden equilibrium — the difference between the open web of agents and the app store, written into law.
- 2 **Model-layer concentration monitoring.** Treat the supply of frontier agents the way supervisors treat systemically important banks: measured market shares, dependency mapping, published concentration telemetry — before the chokepoint forms, not after.
- 3 **Interoperable rails.** Machine-native payments (x402-class and their regulated-currency equivalents), agent-readable public data, and portability of context and contracts — so machine commerce is not forced through proprietary gardens by default.
- 4 **Agent identity and liability law.** Verifiable credentials and a clear human-accountability chain for every transacting agent — the missing layer DeepMind and Microsoft both name, and the institutional franchise awaiting its first mover.

The stewardship framing closes the movement. The machine of Movement II hands society ninety-eight per cent of what its winners create — not out of benevolence, but because someone maintained the competition that forces the donation. The multiplier of Movement III obeys the same rule at higher stakes: **it multiplies whatever structure it is poured into**. Poured into open, competitive, well-governed markets, it multiplies consumer surplus — the Kerala result at civilisational scale. Poured into concentrated, closed, under-governed ones, it multiplies rents and correlated fragility. The technology does not choose; the institutions do.

## THE MOVEMENT IN ONE SENTENCE

Agentic AI is the first technology that upgrades capitalism's own value-creating machinery — and whether society keeps the new ninety-eight per cent is decided not by the models but by **the plumbing: data, identity, payments, and above all competition**.

“

*The most important economic-policy portfolio of the 2030s is stewardship of the giveaway.*

ENSI • THE VALUE ENGINE • MOVEMENT III CLOSE

## MOVEMENT IV — PROOF & ACTION

# The receipts, and the build list

Six countries changed the rules and collected the payout this audit predicts — the same sequence, six different flags. What follows are the receipts; then the nine moves a serious state makes in the next twelve months to collect the agentic one; and the evidence base behind every number in this report, gaps included.

- 37 Country Receipts — China · Poland · Korea · Vietnam · Botswana · Estonia 112
- 38 The ENSI Operating Model — agents, data, plumbing, and the twelve-month matrix 116
- 39 The Evidence Base — 106 documents, 15 angles, and the honest gaps 118

## 37 Country Receipts

The audit's claim is falsifiable: switch on the mechanisms — prices, property, entry, trade, under institutions that hold — and the payout arrives within a generation. Six countries ran the experiment from six different starting points: a communist giant, a post-Soviet plain, a war economy, an ex-colony written off at birth, a middle-income "trap" case, and a rationing-era Baltic state. **All six collected.**

Each receipt below follows the same anatomy — the timeline of the rule change, three audited numbers, what was actually done, and the primary document it traces to. None of these countries was subsidised into wealth, and none got rich on resources alone; the one with the diamonds is the case *for* institutions, not against them. Read the six panels as replications of a single result.

### China — the largest anti-poverty event in history

RECEIPT 01 · 1978–2021

**1978**

Rural reform — farm households regain production decisions

**1980**

First special economic zones open entry and foreign trade

**2001**

WTO accession plugs the workshop economy into world demand

**2021**

Extreme poverty declared eradicated on the national line

~800M

fewer people below the \$1.90 international extreme-poverty line since 1978

World Bank & DRC China, 2022

~75%

of the entire global reduction in extreme poverty over four decades, in one country

World Bank & DRC China, 2022

770M

lifted above China's national poverty threshold since 1978, per the 2021 declaration

World Bank & DRC China, 2022

The reform began where the poverty was: agriculture. From 1978, farm households regained the right to decide what to grow and to sell at the margin; agricultural productivity moved first, exactly as the mechanism catalog predicts. Industrialisation came incrementally — special economic zones from 1980, managed urbanisation and rural-to-urban migration, heavy infrastructure — and WTO accession in 2001 connected the entry boom to world demand. The World Bank's four-decade audit is precise about the recipe: **broad-based economic transformation plus targeted support, carried by effective governance** that coordinated agencies and enlisted non-government actors. Markets did the enriching; state capacity made it stick — the Bank's attribution, and this report's. The largest anti-poverty event in recorded history was not a welfare programme. It was a rule change, run for forty years.

“

*Whether measured with the international or national poverty line, the speed and scale of China's poverty reduction is historically unprecedented.*

WORLD BANK & DRC CHINA · FOUR DECADES OF POVERTY REDUCTION IN CHINA, 2022 · VERIFIED AGAINST THE SOURCE PDF

World Bank & DRC China · Four Decades of Poverty Reduction in China · 2022 · IFI-IGO

World Bank · Poverty & Shared Prosperity 2022

library angle 15 · Case Studies

## Poland — Europe's fastest, from a standing start

RECEIPT 02 · 1989–2013

**1989**

Communism ends; Balcerowicz plan adopted

**1990**

Big-bang: prices freed, currency convertible, trade opened

**2004**

EU accession imports the institutional package wholesale

**2013**

62% of Euro-area-17 income per head — a 500-year record

**62%**

of Western-European (Euro-area 17) GDP per capita by 2013 — more than double the ~28% at transition

Piatkowski, World Bank WPS 6639

**1st**

fastest-growing economy in Europe, 1989–2013, and among the fastest worldwide at its income level

Piatkowski, World Bank WPS 6639

**1500**

the last year Poland's income relative to Western Europe stood this high — before a 300-year decline

Piatkowski, World Bank WPS 6639

Poland entered 1989 with shortages, near-hyperinflation, and a relative-income gap that had been widening since the seventeenth century — from 60–70 percent of the Western level around 1500 to a fraction of it under the plan. The Balcerowicz programme was the region's most compressed rule change: prices freed, budgets hardened, the currency made convertible, trade opened, and private entry unleashed, all at once. It was also the best-paying. Piatkowski's World Bank study names Poland Europe's growth champion of 1989–2013 and dates the arrival of its **"new golden age"**: by 2013, income per head reached 62 percent of the Euro-area benchmark — the country's best position relative to Western Europe since the year 1500. The attribution is to the package, not a policy: liberalisation first, then EU accession importing courts, competition rules, and market access wholesale.

Piatkowski · Poland's New Golden Age · World Bank WPS 6639 · 2013

library angle 15 · Case Studies

## South Korea — the middle-income trap, escaped

RECEIPT 03 · 1962–2001

**1962**

Export-led takeoff — a 29-year run of &gt;6% growth begins

**1980s**

R&amp;D goes private: spending up 26x, past 80% of the total

**1995**

High-income threshold crossed; OECD joined the same year

**2001**

High income regained after the crisis — now innovation-led

**29yrs**

of growth above 6% (1962–1991) — the median rapid-growth run elsewhere lasts nine years

World Bank-KDI, Innovative Korea

**>80%**

of national R&amp;D privately funded by 1990, after a 26-fold rise in private spending through the 1980s

World Bank-KDI, Innovative Korea

**2nd**

highest R&amp;D spending as a share of GDP in the world today — the trap escaped on innovation

World Bank-KDI, Innovative Korea

Korea is the canonical escape from the middle-income trap, and the mechanism of the escape was a handover. The state-directed model that carried the country from low to middle income was deliberately retired in favour of **private-sector-led growth under competition**: through the 1980s private R&D spending grew twenty-six-fold and passed 80 percent of the national total, while policy shifted from directing industries to disciplining them — openness, competition, and a reorientation from adopting foreign technology to generating frontier technology at home. Korea crossed the World Bank's high-income threshold in 1995, was knocked back by the 1998 crisis (–5.1 percent), and returned to high income by 2001, growing 5 percent a year to the global financial crisis. The World Bank-KDI verdict: the trap is escaped by switching engines — from accumulation to private innovation.

World Bank &amp; KDI · Innovative Korea · 2023 · IFI-IGO

library angle 15 · Case Studies

## Vietnam — the China result, replicated small

RECEIPT 04 · 1985–2016

**1985**

GDP per capita US\$231 — a war economy in fiscal crisis

**1986**

Doi Moi "renovation" adopted — a far-reaching reorientation

**2007**

WTO accession, 1 January — trade integration completed

**2016**

GDP per capita US\$2,171 — nine-fold in three decades

**\$231**

GDP per capita in 1985, on the eve of Doi Moi — among the poorest countries in the world

UNU-WIDER WP 2018/126

**\$2,171**

GDP per capita by 2016 — the dragon that rose from the ashes, in UNU-WIDER's title

UNU-WIDER WP 2018/126

**9.4<sub>x</sub>**

rise in income per head over three decades — with poverty falling across regions, not pooling in an enclave

ENSI calculation from UNU-WIDER series

In 1985 Viet Nam's income per head was 231 US dollars. The Doi Moi programme adopted in 1986 was, in UNU-WIDER's assessment, a far-reaching renovation rather than a patch: agriculture decollectivised, prices progressively freed, the dong devalued and unified from its tangle of official rates, foreign investment welcomed, and trade opened step by step — a bilateral agreement with the United States in 2000, WTO accession in 2007, and continued liberalisation since. By 2016 income per head had reached 2,171 dollars, and the study's emphasis falls on **growth with equity**: the transformation raised incomes broadly rather than building an enclave economy. Vietnam is the cleanest small-country replication of the China result — the same sequence, farms then prices then entry then trade, run by a different state at a tenth the scale, with the same direction of payout.

UNU-WIDER · Viet Nam: The Dragon That Rose From the Ashes · WP 2018/126 · UN

library angle 15 · Case Studies

## Botswana — the institutions coefficient, with a flag

RECEIPT 05 · 1966–2001

**Pre-1966**

Tswana assemblies constrain chiefs — property norms predate the state

**1966**

Independence — among the poorest countries on Earth

**1966–2001**

The fastest ~35-year per-capita growth record in the world

**Diamonds**

Resource rents taxed through institutions, not dissolved into them

**~35yrs**

the window over which Botswana out-grew every other country in the world, per capita

Acemoglu, Johnson & Robinson

**1966**

independence, as one of the poorest countries on Earth — the least promising start in the set

Acemoglu, Johnson & Robinson

**~7<sub>x</sub>**

long-run income difference the AJR cross-country work attributes to institutions alone

AJR, AER 2001

Botswana's record — the world's fastest per-capita growth over roughly thirty-five years from an almost destitute start — is usually credited to diamonds, and miscredited: comparable discoveries elsewhere in Africa financed conflict and capture. Acemoglu, Johnson and Robinson attribute the outcome to **institutions of private property**: precolonial Tswana assemblies that constrained chiefs, post-independence elites whose own interests lay in enforcing property rights, and a state that taxed diamond rents through stable, contested institutions instead of being dissolved by them. Botswana is the control group's control group — the same endowment that cursed its neighbours compounded here because the rules were different. In the AJR cross-country estimates, institutions alone account for a roughly seven-fold long-run income difference; Botswana is that coefficient wearing a flag.

Acemoglu, Johnson & Robinson · An African Success Story: Botswana · NBER

AJR · Colonial Origins of Comparative Development · AER 2001

library angle 15 · paraphrase-attributed

## Estonia — liberalisation at machine speed

RECEIPT 06 · 1992–2007

### Jan 1992

Crisis floor: output –30% in a year, inflation over 1,000%

### Mid-1992

Kroon under a currency board — first post-Soviet own currency

### Jan 1994

Flat-rate income tax — later copied by Lithuania, Latvia, Russia

### 2000s

The digital state: e-cabinet, online tax filing, internet banking

# 35<sub>x</sub>

growth in the number of enterprises — from ~2,000 in 1992 to 70,000 by end-1994

Laar, Heritage Bg. 2060

# 89.8%

inflation a year after the currency board — down from over 1,000% in 1992

Laar, Heritage Bg. 2060

# 1<sup>st</sup>

in FDI per capita among transition economies in the late 1990s; an e-government frontrunner since

Laar, Heritage Bg. 2060

January 1992: fuel and food rationed, output down more than 30 percent in a year, inflation above 1,000 percent. Mart Laar's government answered with the region's most radical liberalisation — a balanced budget, the kroon reintroduced under a currency board in mid-1992, tariffs and export restrictions abolished outright (Laar's "free-trade zone"), confiscated property returned to its owners, and, on 1 January 1994, the flat-rate income tax that Lithuania, Latvia and then Russia would copy. Entry answered immediately: **from about 2,000 enterprises in 1992 to 70,000 by the end of 1994**. The same instinct then built the digital state — cabinet business, tax declarations and banking moved online, e-government as lean government. Estonia's receipt is the smallest and fastest in the set: a mid-sized administration rebuilt its economic rails within one political term — precisely the timescale the build list on the next spread assumes.

Laar · The Estonian Economic Miracle · Heritage Backgrounder 2060 · 2007

library angle 15 · Case Studies

## Six receipts, one direction of payout

| CASE        | THE RULE CHANGE                        | MECHANISMS SWITCHED ON                              | HEADLINE RECEIPT                                          |
|-------------|----------------------------------------|-----------------------------------------------------|-----------------------------------------------------------|
| China       | 1978 rural reform → entry → WTO 2001   | Prices, entry, trade — plus state capacity          | ~800M out of extreme poverty; ~75% of the world's decline |
| Poland      | 1989–90 Balcerowicz big-bang           | Prices, competition; EU institutions imported 2004  | 28% → 62% of Western income; Europe's fastest 1989–2013   |
| South Korea | 1980s handover to private R&D          | Innovation under competition and openness           | Middle-income trap escaped; high income 1995              |
| Vietnam     | 1986 Doi Moi renovation                | Prices, entry, trade                                | \$231 → \$2,171 per head, 1985–2016 — with equity         |
| Botswana    | Property-rights institutions from 1966 | Institutions — the rank-0 mechanism                 | Fastest ~35-year per-capita growth in the world           |
| Estonia     | 1992 liberalisation + currency board   | Institutions, trade, entry — then the digital state | 2,000 → 70,000 firms in three years; e-state pioneer      |

### The shared sequence

Free the prices, secure the property, open the entry, open the trade — and hold the institutions steady long enough for compounding to work. No case ran the steps in a different order; no case collected without running them all.

### What the receipts refute

That poverty is destiny (Vietnam, China); that resources are always a curse (Botswana); that reform must be slow to be safe (Estonia, Poland); and that the trap at middle income is inescapable (Korea).

### THE SEVENTH RECEIPT

The agentic receipt is unwritten. The rails it needs — open, legible, contestable markets that agents can read and enter — are a build of Estonian scale, not Chinese scale: one political term. The next two pages are the build list.

## 38 The ENSI Operating Model

What should a state — or any institution of national scale — actually build to capture the multiplier? Three layers: **a portfolio of agents** mapped one-to-one onto the seven mechanisms, **a data architecture** those agents can consume, and **market plumbing** — identity, payments, interoperability — that decides where the surplus flows. Humans own judgement and accountability throughout.

### The agent archetypes — one per mechanism

|                                                                                                                                                                                      |                                                                                                                                                                                    |                                                                                                                                                                                   |                                                                                                                                                                                               |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>PRICES</b><br><b>Transparency agents</b><br>Continuously measure price dispersion and switching friction by sector — the Jensen statistics as live, published national telemetry. | <b>COMPETITION</b><br><b>Consumer reference agents</b><br>Open standards ensuring every citizen's agent can shop, switch and negotiate — keeping the mechanism pointed at markups. | <b>INNOVATION</b><br><b>Research agents</b><br>Run on public compute against open scientific data, attacking the research-productivity decline where private actors under-invest. | <b>ENTREPRENEURSHIP</b><br><b>Founder agents</b><br>The company-in-a-box stack wired into registration, tax and licensing — the falling minimum viable firm size made real, not theoretical.  |
| <b>CAPITAL</b><br><b>Underwriting agents</b><br>In development banking and public procurement — shrinking the Hsieh-Klenow misallocation tax on public capital.                      | <b>TRADE</b><br><b>Export agents</b><br>Package customs, certification and contracting for small firms — commoditising the fixed costs that keep them domestic.                    | <b>INSTITUTIONS</b><br><b>Compliance agents</b><br>Make regulation machine-consumable and enforcement fast and even-handed — compliance as ambient infrastructure.                | <b>THE SANDBOX LAYER</b><br><b>Early-warning agents</b><br>Watch agent transaction volumes, correlation and concentration — as the BIS watches banks. The supervisor's eyes on the new layer. |

#### ABOVE THE PORTFOLIO

Scanning, simulation and red-team agents from the ENSI foresight core sit above the eight archetypes, stress-testing them — a standing adversarial exercise with manipulation, fraud and obfuscation agents, reporting to the same authority that owns the build.

### The data architecture

Agents are only as good as what they can see, and most national data is human-readable at best. The least glamorous layer is the highest-leverage one — it decides whether every mechanism's agents run on facts or on scrapings.

- ▶ **Registries as services** — companies, land, licences, courts exposed as structured, real-time, agent-consumable endpoints.
- ▶ **Cash-flow rails** — e-invoicing and open banking feeding cash-flow-based underwriting.
- ▶ **Machine-format prices** — published price and tariff data in machine formats as a legal default.
- ▶ **Agent-usable science** — scientific and statistical data licensed to permit agent use.

### The market plumbing

Three components decide whether the surplus flows to consumers or to whoever owns the chokepoint:

- ▶ **Identity & credentials** — verifiable agent identity and the who-answers-for-this-agent chain, aligned with the DeepMind trust agenda.
- ▶ **Payments** — machine-native micropayment rails (x402-class and regulated-currency equivalents), so machine commerce is not forced through walled gardens.
- ▶ **Interoperability rights** — a consumer's agent may access any published offer, port any contract, interrogate any service agent, on non-discriminatory terms.

The last is **the competition-policy instrument of the decade** — the difference between the open web of agents and the app store.

Model synthesised from ENSI Report 3, *The Agentic Multiplier* §"The ENSI operating model", drawing on Jensen (QJE 2007) • Hsieh-Klenow (QJE 2009) • Microsoft Research (2025) • DeepMind (2025) • x402 • BIS WP 1194.

## The first twelve months — nine moves, ranked AS OF JULY 2026

The matrix a practitioner should photograph. Ranking is by leverage per unit of political capital: telemetry costs nothing and steers everything; the identity statute is small, boring, and the binding constraint on the entire transactional layer. Preconditions read as a dependency graph — **telemetry first (1)**, **identity before access (4→3)**, **sandbox review before rails (9→5)**.

| # | ACTION                                                                                                                                                                                                                | OWNER                                        | HORIZON | PRECONDITION                                        | METRIC                                                                                                 |
|---|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------|---------|-----------------------------------------------------|--------------------------------------------------------------------------------------------------------|
| 1 | <b>Telemetry index.</b> Stand up the national agentic-economy observatory: Anthropic-Index-style task-usage tracking plus sector price-dispersion series, published quarterly.                                        | Statistics office + AI observatory           | 0–3 mo  | None — existing survey powers suffice               | Index live; dispersion series published for 10 sectors                                                 |
| 2 | <b>Agent-readable markets.</b> Mandate machine-readable prices, terms and switching interfaces in retail finance, energy, telecoms and insurance.                                                                     | Sector regulators                            | 3–12 mo | Format standard written into licence conditions     | Share of offers retrievable by API; median switching time                                              |
| 3 | <b>Open-agent-access rule.</b> Any offer published to humans must be lawfully accessible to their agents on non-discriminatory terms — one clause, aimed at the walled-garden equilibrium while it is still unformed. | Legislature + competition authority          | 6–12 mo | Agent identity framework (move 4)                   | Access-denial complaints resolved; compliance rate at threshold platforms                              |
| 4 | <b>Agent identity &amp; liability law.</b> Verifiable-credential framework and a clear human-accountability chain for transacting agents — minimal statute; the model is Delaware, not Brussels.                      | Justice ministry                             | 6–12 mo | None — the binding constraint today                 | Statute in force; credentialed agents registered                                                       |
| 5 | <b>Micropayment rails.</b> Support x402-class machine-native micropayment and settlement rails, publicly interoperable, with regulated-currency equivalents.                                                          | Central bank + payments regulator            | 6–18 mo | Sandbox-permeability review (move 9)                | Machine-payment volume settling outside walled gardens                                                 |
| 6 | <b>Agent-run compliance as public infrastructure.</b> A free regulatory agent for every one-person firm — incorporation, tax, filings — on three registries made agent-consumable end-to-end.                         | Business registry + tax authority            | 3–12 mo | Registries exposed as structured real-time services | Registration latency, published; one-person-firm formation rate                                        |
| 7 | <b>Research agents into the science system.</b> The R&D budget's next marginal billion into agentic research infrastructure: one pilot on a national scientific strength, public compute, open data.                  | Research ministry + science funders          | 3–9 mo  | Data licences permitting agent use                  | Measured research-productivity effect — the missing Goldman-vs-Acemoglu number, generated domestically |
| 8 | <b>Model-layer concentration monitoring.</b> Track markups and concentration at the model layer as seriously as bank capital — the 2.2% split is defended here or lost here.                                          | Competition authority + financial supervisor | 0–6 mo  | Telemetry index (move 1)                            | Model-layer concentration and markup series, published                                                 |
| 9 | <b>Intentional sandbox.</b> Sandbox the agent economy on purpose — DeepMind's intentional-and-controlled quadrant — before it entangles with human markets by accident.                                               | Central bank + financial supervisor          | 6–12 mo | Early-warning agents deployed                       | Agent transaction volume, correlation and concentration dashboards live                                |

### WHAT DECIDES THE SPLIT

Moves 3, 5 and 8 are the competition instruments. They — not the technology — decide whether the producers of intelligence keep 2 percent of the value this time, or 40.

### DONE, BY JULY 2027, LOOKS LIKE

A quarterly telemetry report the cabinet reads; three registries an agent can use end-to-end, latency published; the identity statute and access rule in force; one research-agent pilot reporting measured productivity effects.

Owners and horizons are ENSI recommendations (ENSI estimate), synthesised from Report 3 §"The first twelve months, ranked — as of July 2026". **Anthropic** • Economic Index **DeepMind** • Virtual Agent Economies • 2025

**x402** • whitepaper

**BIS** • WP 1194

**Microsoft Research** • The Agentic Economy • 2025

## 39 The Evidence Base


Every number in this report traces to a **downloaded primary document** — **106 of them, across 15 research angles** (~287 MB, compiled 8 July 2026), each pulled from the original institution or a verified open-access mirror and catalogued with its provider, type, source URL and file name in its angle's own index. Quotes were verified against the source PDFs before use; back-of-envelope aggregations are labelled ENSI estimates.

### The fifteen angles

| #  | ANGLE                                      | DOCS | ANCHOR SOURCES                                                  |
|----|--------------------------------------------|------|-----------------------------------------------------------------|
| 01 | Foundations — What Capitalism Is           | 8    | Hayek (AER 1945) · Friedman · Besley-Ghatak                     |
| 02 | Who Captures the Value — Consumer Surplus  | 7    | Nordhaus (NBER w10433) · Brynjolfsson-Collis-Eggers (PNAS 2019) |
| 03 | The Great Enrichment — Growth & Poverty    | 7    | Maddison Project · World Bank PSP 2022 · Deaton Nobel lecture   |
| 04 | Innovation & Creative Destruction          | 7    | Jones-Summers (NBER w27863) · Bloom et al. (AER 2020)           |
| 05 | Entrepreneurship & Business Dynamism       | 6    | Haltiwanger-Jarmin-Miranda (w16300) · Decker et al.             |
| 06 | Competition & the Consumer                 | 7    | Borenstein-Rose · De Loecker-Eeckhout (w23687) · CMA 2024       |
| 07 | Capital Allocation & Financial Markets     | 7    | Levine · Lerner-Nanda (HBS) · Hsieh-Klenow (QJE 2009)           |
| 08 | Trade, Globalization & Value Chains        | 7    | Hufbauer-Lu (PIIE) · Frankel-Romer · WDR 2020                   |
| 09 | Institutions, Property Rights & Freedom    | 8    | AJR (AER 2001) · Fraser EFW 2024 Ch.1 · B-READY 2024            |
| 10 | Human Development Dividends                | 7    | Preston (1975) · Cutler-Deaton-Lleras-Muney · UNDP HDR          |
| 11 | Critiques Answered — Failures & Correction | 7    | WIL 2022 · Autor-Dorn-Hanson · Ostrom Nobel lecture             |
| 12 | AI & the Productivity Frontier             | 7    | Goldman Sachs 2023 · Acemoglu 2024 · NBER w31161                |
| 13 | The Agentic Engine — Agents as Actors      | 7    | Anthropic EI · DeepMind 2025 · x402 · BIS WP 1194               |
| 14 | Measuring Value — Beyond GDP               | 7    | Jones-Klenow (AER) · GDP-B (w25695) · Coyle                     |
| 15 | Case Studies — Countries & Firms           | 7    | World Bank Four Decades · Jensen (QJE 2007) · Laar              |

### Where the corpus comes from

FIG. E-1 • SOURCE MIX, BY DOCUMENT TYPE



More than half the corpus is peer-reviewed or working-paper economics (56 of 106); official institutions supply another 33. Industry sources (Anthropic, Goldman Sachs, Microsoft, DeepMind, x402) carry the agentic frontier only, never the historical claims.

## The marquee sources

The load-bearing documents, grouped by the job they do in the argument. Every chip is a file on disk; a handful of anchor papers are deliberately indexed in two angles (Nordhaus in 02 and 11, GDP-B in 02 and 14, Generative AI at Work in 12 and 13, the Deaton lecture in 03 and 10) because they carry both arguments.

### THE MECHANISM — WHY MARKETS CREATE VALUE

**Hayek** · Use of Knowledge in Society · AER 1945   **Nordhaus** · Schumpeterian Profits · NBER w10433  
**Fraser** · EFW 2024 Ch.1   **AJR** · Colonial Origins · AER 2001   **De Loecker-Eeckhout** · Markups · w23687  
**Hsieh-Klenow** · Misallocation · QJE 2009   **Lerner-Nanda** · VC's Role · HBS   **Kortum-Lerner** · VC & Innovation  
**CEA** · Competition brief · 2016

### THE EVIDENCE — HOW MUCH, FOR WHOM

**Maddison Project** · long-run GDP series   **World Bank** · Poverty & Shared Prosperity 2022  
**Jones-Summers** · Calculus of Innovation · w27863   **Bloom et al.** · Ideas Harder to Find · AER 2020  
**Haltiwanger-Jarmin-Miranda** · w16300   **Hufbauer-Lu** · Payoff to Globalization · PIIE   **Deaton** · Nobel lecture

### THE HONEST LEDGER — DEBITS AND MISMEASUREMENT

**WIL** · World Inequality Report 2022   **Autor-Dorn-Hanson** · China Shock   **Ostrom** · Nobel lecture  
**Brynjolfsson-Collis-Eggers** · PNAS 2019   **Brynjolfsson et al.** · GDP-B · w25695  
**Jones-Klenow** · Beyond GDP · AER

### THE AMPLIFIER — AI AND AGENTS

**Goldman Sachs** · Global Economics · 2023   **Acemoglu** · Simple Macroeconomics of AI · 2024  
**Brynjolfsson-Li-Raymond** · GenAI at Work · w31161   **Noy-Zhang** · Science 2023  
**Eloundou et al.** · GPTs are GPTs · 2023   **Anthropic** · Economic Index  
**DeepMind** · Virtual Agent Economies · 2025   **Microsoft Research** · The Agentic Economy · 2025  
**x402** · whitepaper   **BIS** · WP 1194

### THE PROOF — CASES

**World Bank & DRC** · Four Decades (China)   **Piatkowski** · Poland · WPS 6639  
**World Bank-KDI** · Innovative Korea   **UNU-WIDER** · Viet Nam · 2018/126   **AJR** · Botswana · NBER  
**Laar** · Estonian Miracle · Heritage 2060   **Jensen** · The Digital Provide · QJE 2007

## Missing & substituted — the honest gaps

Six sources hard-blocked scripted download or proved unusable. Each gap and its substitute is recorded in the relevant angle's index; none is load-bearing alone.

- ▶ **Fraser EFW 2024, full report** — fraserinstitute.org 403s non-browser clients. Used instead: the **Cato-hosted Chapter 1**, the identical core index chapter.
- ▶ **IMF Staff Discussion Notes** — imf.org blocks scripted requests. SDN 15/08 obtained via the **IMF eLibrary endpoint**; SDN 15/13 and SDN/2024/001 remain missing, no reliable mirror.
- ▶ **McKinsey, Economic Potential of Generative AI** — media endpoint bot-blocked, no stable mirror. Macro range covered by **Goldman Sachs and OECD**.
- ▶ **Gornall-Strebulaev, VC evidence** — SSRN and Stanford GSB bot-gated; the NBER-numbered URL served a **different paper and was removed**. Ground covered by **Lerner-Nanda**.
- ▶ **WHO, Macroeconomics and Health (CMH 2001)** — iris.who.int serves a JS interstitial, not the PDF. Ground covered by **Preston and the Deaton papers**.
- ▶ **Google, Agent Payments Protocol (AP2)** — published as web/GitHub content only. Used instead: the **x402 whitepaper**, the machine-payments primary source.

### THE STANDARD THIS BUYS

Every claim in this report can be checked against a file on disk — provider, URL, and download status recorded per document. Where a number is an ENSI aggregation rather than a source figure, it says so on the page where it appears.

# The correct posture toward the agent economy is stewardship: measure it, open it, and defend the 98 percent giveaway.

Agents in competitive markets are surplus machines for everyone; agents in concentrated markets are the most efficient rent collectors ever deployed. Nothing in the technology chooses between those futures. Institutions choose.