

The Asymmetry Engine, Chapter 15: Atoms, hands and inertia

Martin Roark

Chapter 15 of *The Asymmetry Engine: How AI Closes the Gaps, and What Stays Scarce*, October 2026. The online edition, with the interactive figures and the links between chapters, is at <https://future-seems-so-good.com/blog/the-asymmetry-engine#inertia>. License: [CC BY 4.0](#).

15:1 Abundant intelligence changes the world at the speed of the slowest step in each chain that closes a gap, and the slowest steps are made of atoms, permits and people. Where a gap has a physical part that intelligence cannot substitute, its closing rate converges to the rate at which that part is supplied, while the inertia measured against intelligence grows with intelligence. Inertia therefore sets the timing of the closing more than its destination. A top-down protocol can break inertia by enabling bottom-up competition, as Brazil's Pix did for payments, and what only people can do grows dearer as what machines make grows cheap.

15.1 Why there is so much inertia

15:2 Inertia is the term φ_k that divides every closing rate in the master equation, $\lambda_k = I_k v_k a_k / \varphi_k$. It has five sources: sunk physical capital, know-how that must be learned in place, accountability and regulation, coordination among many parties, and incumbents' rents that depend on the inertia itself. None is new. What is new is that the numerator, the intelligence supply, is about to grow by orders of magnitude, so the question becomes which term binds.

15:3 The canonical measurement is Paul David's. In 1899 electric motors supplied "less than 5 percent of factory mechanical drive", and reaching half took "another two decades, roughly speaking". Electrification had no effect on manufacturing productivity "before the early 1920s ... four decades after the first central power station opened for business" (David 1990). Serviceable steam-era plants were not worth replacing; early adopters laid electric "group drive" over the old shafts and belts instead of redesigning the factory; momentum came only after 1914–17, when regulated utility rates fell; and redesign needed "a cadre of experienced factory architects and electrical engineers", built by a learning process "inherently uncertain and slow to gain momentum". When the payoff came, motor capacity accounted statistically for about half of the five-point acceleration in US manufacturing productivity growth in 1919–29.

15:4 The lag generalizes. Across 15 technologies and 166 countries, adoption came on average 45 years after invention, with a standard deviation of 39 years, and newer technologies were adopted faster (Comin and Hobijn 2010). Brynjolfsson, Rock and Syverson call implementation lags "likely ... the biggest contributor" to the modern productivity paradox (2017), because general-purpose technologies "enable and require significant complementary investments", mostly intangible and badly measured; counting them puts US total factor productivity 15.9% above the official measure by the end of 2017 (2021). The lag is the world reorganizing around the invention.

15:5 The 2025–26 data show two diffusions at different speeds. The tool spreads fast: three years after ChatGPT's launch, 54.6% of US adults aged 18 to 64 used generative AI, against 19.7% for the personal computer and 30.1%

for the internet at the same age (St. Louis Fed). The reorganization spreads slowly, as the firm and executive surveys of §8.1 record. The counter-evidence is young: US productivity grew 2.8% in 2025 (BLS), which Brynjolfsson reads as the start of a "harvest phase", and the first two quarters of 2026 were softer. The widely quoted "95% of AI pilots fail" misreads a small report whose own funnel implies that a quarter of pilots reached production (MIT NANDA).

15:6

Definition 15.1 (Diffusion with inertia; the J-curve). Let a technology's adopted share $s_{\text{ad}}(t)$ follow Bass's law with innovation and imitation coefficients k_{inn} and k_{imi} , slowed by inertia φ :

$$\dot{s}_{\text{ad}} = \frac{(k_{\text{inn}} + k_{\text{imi}} s_{\text{ad}})(1 - s_{\text{ad}})}{\varphi}$$

so every adoption date stretches by the factor φ . When the gains need complementary investment that takes time to build and is expensed as it is made, measured productivity first dips and then rises after a lag: the **productivity J-curve**.

▼ *The closed form and the measured dip*

The solution is $s_{\text{ad}}(t) = \frac{1 - e^{-(k_{\text{inn}} + k_{\text{imi}})t/\varphi}}{1 + (k_{\text{imi}}/k_{\text{inn}}) e^{-(k_{\text{inn}} + k_{\text{imi}})t/\varphi}}$, and adoption is fastest at $t_{\text{pk}} = \varphi \ln(k_{\text{imi}}/k_{\text{inn}})/(k_{\text{inn}} + k_{\text{imi}})$. If the gains need complements that take T_{cmp} to build and cost c_{cmp} per unit of adoption, expensed and unmeasured, measured productivity is

$$y_{\text{meas}}(t) = k_{\text{gain}} s_{\text{ad}}(t - T_{\text{cmp}}) - c_{\text{cmp}} \dot{s}_{\text{ad}}(t)$$

with k_{gain} the eventual gain. It is negative while adoption is fast and the complements immature, and positive once they mature.

15:7 Calibrate Definition 15.1 with the meta-analytic averages $k_{\text{inn}} = 0.03$ and $k_{\text{imi}} = 0.38$ (Sultan, Farley, and Lehmann 1990). With $\varphi = 1$, adoption goes from 5% to 50% in about five years and is fastest after about six. Factory electric drive took about twenty years for the same leg (15:3), an inertia of $\varphi \approx 4$, and the lag of the complements brings the total to David's four decades. Generative AI as a personal tool passed 54% of US adults aged 18 to 64 in three years (15:5), which implies $\varphi \approx 0.4$. Read the share of US work hours that it assists, 6.3% after about three and a half years, as an adoption share, and the reorganization has $\varphi \approx 2$. The two inertias differ about fivefold, and the J-curve says the measured payoff follows the slower one.

15.2 The weakest link

15:8 Aghion, Jones and Jones show that growth can be constrained "by what is essential and yet hard to improve", and that such a constraint can bind "even with complete automation and even with a superintelligence" (2017). Cowen gives the institutional case, in which the rate of drug approval sets the pace of medicine (§14.3). Inside AI research the same form makes compute the homeostat (Proposition 11.2). For the economy as a whole, the physical world plays that part.

15:9 Split the wall-clock time of a process into a cognitive part (design, code, analysis, review) and a physical part (fabrication, construction, permitting, shipping, a culture growing, a trial enrolling). The **physical share** s_{phys} is the physical part's fraction of the whole. If cognition becomes k_{cog} times faster, the process speeds up by

15:10

$$\text{speedup}(k_{\text{cog}}) = \frac{1}{s_{\text{phys}} + (1 - s_{\text{phys}})/k_{\text{cog}}} \leq \frac{1}{s_{\text{phys}}} \quad (15.1)$$

15:11 This is Amdahl's law applied to matter. A process that is half physical never runs more than twice as fast; one that is a tenth physical gains 5.3× from tenfold faster cognition and at most 10× from any cognition. In knowledge work the physical part is often human review, which takes 0.27 of a GDPval task's time and 0.24 of its cost and sets the ceilings of (2.3). A digital twin lowers the physical share only inside its validated envelope (Definition 14.2), and a research loop under a moving physical cap grows at the slower of its own speed and the build-out's (Proposition 11.2).

15:12 Long chains add a second limit, because their reliabilities multiply.

15:13

Definition 15.2 (The O-ring product). In the **O-ring product** of [Kremer \(1993\)](#), expected output is proportional to the product of the probabilities that each task is done right. A product with N_{task} independent features, each working with probability p_{task} , gives a session without a defect with probability $p_{\text{task}}^{N_{\text{task}}}$; at 1,000 features and $p_{\text{task}} = 0.999$ that is $e^{-1} \approx 37\%$. If each feature's defect probability falls tenfold for every h_{nine} hours of effort spent on it, from k_{def} at none, a product reliability R_{tgt} needs $1 - p_{\text{task}} \approx (1 - R_{\text{tgt}})/N_{\text{task}}$ on every feature and a total effort of

$$N_{\text{task}} h_{\text{nine}} \log_{10} \frac{k_{\text{def}} N_{\text{task}}}{1 - R_{\text{tgt}}}$$

which grows like $N_{\text{task}} \ln N_{\text{task}}$. Each extra "nine" costs a constant effort per feature, a "[march of nines](#)" in Andrej Karpathy's phrase, and lifting the example from 0.37 to 0.99 takes two more nines on all 1,000 features.

15:14 The O-ring product explains why products worth a trillion dollars still fail at their margins: a very large N_{task} meets a fixed budget of human effort, which people find aversive (§6.1), so the $N \ln N$ bill is never fully paid. Such products exist only because thousands of people coordinate through hierarchies, specifications and markets, far beyond a natural group size of about 150 ([Dunbar 1992](#)), and by Conway's law they inherit their makers' coordination losses. The institutions that let brains selected for small bands do this are what makes organizations inert (§8.1). Machine effort carries no aversion, scales with gigawatts and pays the bill first where checking is cheap, in software (Definition 5.3).

15:15 Closing a real gap combines cognitive work with physical work (labor, permits, materials, machines), and the two are complements.

15:16

Proposition 15.1 (The binding link). Let closing gap k combine cognitive work, supplied at $\nu_{\text{cog}} = I_k v_k a_k / \varphi_0$ with φ_0 the inertia of purely digital work, and physical work supplied at ν_{phys} , through a CES aggregate with weight s_{cog} on cognition and elasticity of substitution $\varepsilon_s < 1$:

$$\lambda_k = \left[s_{\text{cog}} \nu_{\text{cog}}^{\frac{\varepsilon_s - 1}{\varepsilon_s}} + (1 - s_{\text{cog}}) \nu_{\text{phys}}^{\frac{\varepsilon_s - 1}{\varepsilon_s}} \right]^{\frac{\varepsilon_s}{\varepsilon_s - 1}}$$

Define the effective inertia $\varphi_k = I_k v_k a_k / \lambda_k$, the value that makes (0.2) hold exactly. Then, as $I_k \rightarrow \infty$:
 (i) $\lambda_k \rightarrow (1 - s_{\text{cog}})^{\varepsilon_s / (\varepsilon_s - 1)} \nu_{\text{phys}}$, so physical supply bounds the closing rate; (ii) the intelligence elasticity of the closing rate, $\partial \ln \lambda_k / \partial \ln I_k$, tends to zero; (iii) the effective inertia grows in proportion to I_k , so inertia is relative.

▼ *Proof*

With $\varepsilon_s < 1$ the exponent $(\varepsilon_s - 1) / \varepsilon_s$ is negative, so the cognitive term vanishes as $\nu_{\text{cog}} \propto I_k$ grows: (i). The elasticity is the cognitive term's share of the bracket, $\left[1 + \frac{1 - s_{\text{cog}}}{s_{\text{cog}}} (\nu_{\text{cog}} / \nu_{\text{phys}})^{(1 - \varepsilon_s) / \varepsilon_s} \right]^{-1} \rightarrow 0$: (ii). Substituting (i) into $\varphi_k = \varphi_0 \nu_{\text{cog}} / \lambda_k$ gives $\varphi_k \simeq \varphi_0 \nu_{\text{cog}} / ((1 - s_{\text{cog}})^{\varepsilon_s / (\varepsilon_s - 1)} \nu_{\text{phys}})$, linear in I_k : (iii). ■

15:17 A worked case gives the size. With $\varepsilon_s = 0.5$, equal weights and the two supplies equal at the start, a hundredfold rise in intelligence raises the closing rate 1.98-fold and the measured inertia about 50-fold. This is relative inertia in Hannan and Freeman's sense, where structures are inert "when the speed of reorganization is much lower than the rate at which environmental conditions change" (1984): the same physical world grows more inert as intelligence quickens. Where the closing rate stops answering to intelligence, the rent moves to the physical link, since only binding inputs earn rent ((16.1)); that is why the build-out's margins pool in memory, turbines and electrical labor (§16.3). The physical world is the outer brake around every loop that must act on matter, the research loop of Figure 11.1 included.

15:18 The strongest rebuttal comes from Erdil and Besiroglu. They find Baumol-type arguments "relatively weak in light of the massive transitory gains in output" that such constraints still permit, show that steady automation of complementary tasks gives growth that is "initially slow" and "extremely fast towards the end", and judge that "delays on the order of 30-40 years seem like the slowest that takeoff could end up being" (2023). Their most plausible slow case is this chapter's subject, "physically embodied" tasks such as general-purpose robotics", and if robots come to build the fabs and power plants, the cap joins the loop (15:32). Roy Amara's law was about diffusion from the start: "we consistently overestimate the rate of diffusion and the impacts of technology in the short run but underestimate diffusion and impact in the long run" (Quote Investigator). The binding link predicts effects that are delayed and back-loaded, and they can still be large.

15.3 Connecting things

15:19 Intelligence acts on the physical world only through what is connected to it. Firm AI adoption "likely depends on the extent to which existing workflows are digitized", and fewer than 10% of businesses in agriculture, transportation, accommodation and food services, and construction report using AI (Minneapolis Fed). Connection and control are also different things. About 21.1 billion IoT devices were connected at the end of 2025, and fewer than 1% had a true edge-AI component (IoT Analytics); a device that only reports cannot act.

Proposition 15.2 (The actuation bound). Let a regulator act on a physical system through an actuator set R_{act} and observe it through sensor readings M_{sens} . Whatever its intelligence,

$$H(O) \geq H(D) - \log_2 |R_{\text{act}}|, \quad \Delta H_{\text{closed}} \leq \Delta H_{\text{open}} + I(D; M_{\text{sens}})$$

where ΔH_{closed} and ΔH_{open} are the entropy it removes with and without feedback. An unconnected system has $|R_{\text{act}}| = 1$, so $H(O) \geq H(D)$ and intelligence removes nothing; with actuators and no sensing, $I(D; M_{\text{sens}}) = 0$ and only open-loop control is possible. Machines that are not connected cannot be regulated.

▼ *Proof*

The first inequality is requisite variety, (4.1), with the response confined to the actuators. The second is Touchette and Lloyd's bound on what feedback adds (4:9). ■

15:21 Intelligence multiplies the value of a control channel and cannot create one. Sensors, interfaces, actuators and the wiring of operational machinery into information systems are therefore complements of intelligence in the sense of Proposition 15.1, and their supply is the ν_{phys} of many gaps. Above them sit the hands that wire models into live systems (§3.4), and above those the task Benedict Evans names: "The hard part is knowing that you need a tool for this in the first place, and then knowing what the tool should do" (2026). Until robots are general, the most general actuators an AI can reach are human hands.

15:22 Connecting a power plant is itself a queue. The median US power project completed in 2024 took 55 months from interconnection request to commercial operation, against 36 months in 2015 and 22 in 2008; about 2,290 GW sat in the queues at the end of 2024, and only 13% of the capacity that requested interconnection from 2000 to 2019 had been built (LBNL, *Queued Up*). Federal environmental impact statements took 4.5 years on average from 2010 to 2018, and a quarter took more than six (CEQ). Turbines and transformers carry lead times of years (§16.1), and medicine runs on trial clocks of a decade (§14.4).

15:23 Some of this inertia is deliberate. A permit or a trial is a loop built to be slower than what it regulates, the homeostasis of §6.5, held against the danger Wiener named in 1960, a machine too fast to interrupt once it has started (4:56). A delayed brake can hold only a loop that grows by less than a factor of e within the brake's lag (Proposition 4.3). The policy question is which inertia protects and which is merely inherited.

Forecast 15.1 (Inertia rises in relative terms). By the end of 2027, the latest edition of *Queued Up* still puts the median time from interconnection request to commercial operation for US power projects above four years, while frontier capability keeps compounding, so physical inertia grows relative to intelligence as Proposition 15.1 predicts.

Horizon 2027-12-31

Probability 85%

Check Read the median duration for projects completed in the latest year covered by the latest edition of LBNL's *Queued Up* published by the horizon. A median of four years or less falsifies the claim.

15.4 Hands and robots

15:25 Demand for physical work arrives years before robots can meet it. US construction needs about 349,000 net new workers in 2026 and 456,000 in 2027, with data centers the main driver ([ABC](#)); 81% of contractors with electrician openings struggle to fill them ([AGC](#)); and CSIS models 63,000 to more than 140,000 extra skilled-trades workers for AI infrastructure by 2030 ([CSIS](#)). Between 45% and 70% of a data center's construction budget goes to electrical subcontractors, and an Abilene contractor who can pay \$20 an hour lost workers to sites paying \$35 plus overtime ([Texas Tribune](#)). "Electricians cannot be trained overnight", and apprenticeships must expand by half by 2030 ([CSIS](#)): an apprenticeship is a build time.

15:26 The demand is front-loaded and uneven. A large data center "can employ up to 1,500 construction workers" and keeps about 50 full-time staff once built ([Fortune](#)). Physical automation already displaces hiring without humanoids: internal Amazon documents envisage avoiding more than 600,000 US hires by 2033 through warehouse robotics ([New York Times](#)). The claim holds for the build-out in electrical and mechanical trades and in robot-data work, and it is weak in warehousing and logistics, where non-humanoid automation is already ahead.

15:27 The hands are already training their replacements. Humanoid makers pay people \$20–30 an hour to teleoperate robots or to wear sensor suits that generate training data ([Figure](#), [Dexmate](#)). A binding constraint's rent is also the bounty for relaxing it, so the first things worth automating are the binding constraints themselves.

15:28

Forecast 15.2 (The electrician premium). Through 2028, wages in US electrical trades keep growing faster than the private-sector average while data-center construction spending rises, so the scarcity at the physical link of Proposition 15.1 shows up as a premium.

Horizon 2028-12-31

Probability 60%

Check Compare growth in BLS average hourly earnings for electrical contractors and other wiring installation contractors with growth in private-sector average hourly earnings, from December 2025 to the latest month published by the horizon, and check Census construction spending on data centers over the same period. The forecast holds if electrical earnings grew faster while that spending rose; it fails otherwise.

15:29 Robots are coming, and in 2026 their scale is small. Counts of humanoids shipped in 2025 run from about 18,000 ([IDC](#)) through about 13,000, some 90% of them Chinese ([Omdia](#)), to nearly 7,000 on the [IFR](#)'s narrower definition, and more than 85% of 2025 deployments were performances, education, data collection or tours ([IDC](#)). [Counterpoint](#) counts more than 22,000 in the first half of 2026. Factories, by contrast, installed 603,307 industrial robots in 2025 and operated 5,079,078 ([IFR](#)). Goldman Sachs forecasts humanoid shipments of 75,000 in 2026, 890,000 in 2030 and 6.5 million in 2035 ([24/7 Wall St](#)). Even at the 2030 rate in every year from 2026, the 2030 stock would stay under 4.5 million, below the industrial fleet of 2025, and 75,000 humanoids for the world in 2026 compare with 349,000 missing construction workers in the United States alone.

15:30 The binding constraint in robotics is almost literally the hand. Every Optimus volume target has slipped. Tesla missed its 2025 goal of thousands of useful robots and reportedly stopped mass production that year to redesign the hands and forearms ([heise](#)); in January 2026 Musk said Optimus was "still in the R&D phase ... not in usage in our factories in a material way" ([Electrek](#)). Tesla ended Model S and X production to free Fremont for

Optimus lines whose first builds collect training data ([Tesla](#)). By August they reportedly made several hundred a week, all used internally ([Electrek](#)), and as of October 2026 Tesla had disclosed no external Optimus customer or revenue. In July Musk was more cautious than ever: "those demonstrations you're seeing are preprogrammed or remote controlled. So there is no humanoid robot that is actually able to do generalized tasks" ([earnings call](#)). The "1 Million Bots Delivered" target is a milestone in his pay package ([SEC](#)). The best-documented deployment is modest: Figure's robots spent 11 months at BMW's Spartanburg plant, logging more than 1,250 operating hours and loading more than 90,000 parts, with no public intervention rates ([Figure](#)), and Hyundai plans Atlas for parts sequencing from 2028 and assembly from 2030 ([Boston Dynamics](#)). I read the industry's effort as real, its direction as plausible and its dates as unreliable.

15:31 The labor asymmetry behind the bet is large and visible. At Goldman's projected 2030 average price of about [\\$30,800](#) and an assumed 20,000 operating hours of life, a humanoid's capital costs about \$1.54 an hour before power, upkeep and supervision, against the \$35 an hour of a data-center electrician in Texas (15:25). An asymmetry everyone can see gets priced, and Musk has said that "[~80% of Tesla's value will be Optimus](#)". The cost curve is Chinese: Unitree's average humanoid price fell 72% in two years to about \$25,000, the firm reported a 2025 profit, and its shares rose 460% on their first trading day, 19 August 2026 ([TechNode](#)), while China's planner warns of more than 150 humanoid makers and "redundant products" ([NDRC](#)). If humanoid labor becomes cheap and widely supplied, competition passes the labor-cost gap to buyers (Proposition 16.3), and producers keep the rents of the inputs that bind: actuators, reducers, hands, inference chips and teleoperation data. Nothing here is investment advice.

15:32 Moravec's paradox is the usual explanation for the lag, and it has never been empirically tested ([Narayanan](#)). The physical share gives a firmer one. A robot is matter: it must be designed, certified, manufactured, installed and maintained, and it joins an installed base that turns over across decades-long asset lives, whatever the state of the algorithms. Robotics is also what would let the physical base grow with intelligence. Until robots build the fabs and the power plants, hands are the coupling between cognition and matter, and they are paid accordingly.

15.5 A protocol against inertia

15:33 Brazil's banks put numbers on inertia. The four largest held 54.7% of assets, 57.1% of deposits and 57.9% of credit in 2024 ([Valor](#)). The lending spread, 32.5 percentage points in 2024, is among the highest in the world, and the central bank attributes most of it to defaults, overhead and taxes, with bank margin about a fifth ([BCB](#)). Why a bank moves more slowly than a mobile game is the subject of §5.8.

15:34 Pix broke the inertia of payments. The central bank built the instant-payment system and operates it, made participation mandatory for every institution with more than 500,000 active accounts ([Resolução BCB nº 1](#), 12 August 2020), and fixed its service level: a median payer experience of 6.0 seconds and a 99th percentile of 10.0 seconds, around the clock ([BCB](#)). Live from 16 November 2020, Pix handled 79.8 billion transactions worth more than R\$35 trillion in 2025, and about 86% of adults use it ([BCB](#)). The next-day DOC transfer, created in 1985, was abolished in 2024 ([Febraban](#)). Competition followed: a doubling of Pix value is associated with a 14-basis-point narrowing of the deposit-rate gap between small and large banks ([Sarkisyan](#)). A mandate from the top moved an industry in five years where a weak reward signal alone had not.

15:35 Pix is a top-down protocol that enables bottom-up competition. The incumbents' goals were misaligned with the public's, since their rents depended on float, fees and slow transfers, and in the topology model a group with low alignment ($a = 0.2$) does best fully top-down at every level of intelligence (Experiment 7.1; Proposition 7.1). A regulator with a clear objective imposed one protocol, and thousands of institutions and apps then competed on top of it, as markets and hives do (§7.7). Inertia born of coordination failure and incumbent rents can be broken fast this way, because the protocol is digital and the inertia it removes is institutional.

15:36 What five years of Pix did not move is the profit pool, which rests on underwriting, relationships and reputation. Nubank, founded in 2013, is the kind of entrant through which most organizational change arrives ([Hannan and Freeman](#)), and it serves about 62% of Brazilian adults ([Nubank](#)) at a monthly cost to serve of about \$1 per active customer ([Nu Holdings](#)), yet its management puts its share of the banking profit pool near 7%, with incumbents' average revenue per active customer at \$40–45 against its \$17 ([earnings call](#)). Speed and design win users quickly, and inertia keeps the money for years. That is the shape to expect when cheap intelligence meets incumbents: the interface changes fast and the rents slowly.

15:37 The pattern makes a prediction: inert sectors transform fastest through mandated open protocols with private competition on top, such as open finance, health-data interoperability and digital permitting, and slowest through incumbents' internal AI programs.

15:38

Forecast 15.3 (Protocols beat programs). By the end of 2030, countries with open public protocols (instant payments, digital identity, open finance) show larger AI-attributable productivity gains in banking or health than comparable countries without them.

Horizon 2030-12-31

Probability 35%

Check Read the cross-country studies published through 2030 by the OECD, IMF, BIS or World Bank, or in peer-reviewed journals, that attribute productivity or cost-to-serve gains in banking or health to AI. The forecast holds if most of those that split countries by whether such protocols exist find the larger gains where they do; it fails if most find the larger gains where they do not or no difference, or if none makes the split.

15.6 The price of things

15:39 Cheap intelligence lowers the price of what it makes and raises the relative price of what it cannot make.

15:40

Proposition 15.3 (Baumol's two sectors). Let automatable goods have productivity growing at g_g and human-intensive services at $g_s < g_g$, with one wage w across both (Baumol 1967). With prices equal to unit labor costs, the price of services p_s relative to the price of goods p_g rises at the gap in productivity growth,

$$\frac{d}{dt} \ln \frac{p_s}{p_g} = g_g - g_s > 0$$

while the wage-price of goods, p_g/w , falls at g_g . As machines make things cheap, what only hands can do grows dearer.

15:41 The US record fits. From 2000 to 2025 the consumer price index for televisions fell 98%, computers 92%, toys 74% and software 73%, while hospital services rose 274%, college tuition 188% and day care 147%, against 87% for all items ([BLS indexes, compiled by Chartive](#)). Anthropic's own illustration of AI's effect is Baumol's: "Teachers may prepare lesson plans more efficiently with AI while having no impact on time spent with students in the classroom" ([Anthropic](#)).

15:42 The proposition splits "the price of things will go down" into two predictions. Whatever AI and robots can make gets cheaper in wage terms, as goods have for a quarter-century. Whatever stays human-intensive gets relatively dearer, and where demand for it is inelastic its share of spending rises and aggregate growth converges toward its slow rate. Robots are the solvent, because they move tasks from the slow sector to the fast one and raise g_s , at a speed set by the inertia of Definition 15.1 and the physical link of Proposition 15.1; falling robot prices then expand robot use through the Jevons effect ((3.3)). They take structured tasks such as cleaning, logistics and food preparation first, and care later. Cheap intelligence and cheap hands together are the one combination that dissolves Baumol's cost disease.

15:43 AI itself already shows both sectors. The price of a fixed capability falls steeply (3:48), while its physical inputs grow dearer: conventional DRAM contract prices are forecast up 10–15% quarter on quarter in the fourth quarter of 2026 ([TrendForce](#)), 2027 HBM prices up 121% ([TrendForce](#)), and TSMC raises prices by up to 10% from 2027 ([Nikkei](#)). Altman wrote the same split in early 2025: "the price of luxury goods and a few inherently limited resources like land may rise even more dramatically" ([Three Observations](#)). The price of things falls for things made mostly of cognition and rises for things made of memory, turbines and land; §16.3 follows those rents, and §17.1 asks who bears the transition.

15:44

Forecast 15.4 (The Baumol inversion in physical services). By the end of 2035, humanoids generalize and the relative price of labor-intensive physical services stops rising: world humanoid shipments pass one million a year, and day care, home care and cleaning no longer outpace the all-items CPI by more than a percentage point a year.

Horizon 2036-06-30

Probability 30%

Check Compare the average yearly change from 2031 to 2035 in the CPI indexes for day care and preschool, home health care and, where BLS publishes it, domestic services with that of the all-items CPI, and read world humanoid shipments from IDC or the IFR. The forecast holds if shipments passed one million in at least one year from 2031 to 2035 and those services outpaced the all-items CPI by one percentage point a year or less; it fails otherwise.

15:45 Inertia sets the timing more than the destination, and two 2026 models bracket the range. In one, automating 13% of tasks across all sectors suffices to push the economy into explosive growth ([Davidson, Halperin, Houlden and Korinek](#)). In the other, automation that continues historical patterns lifts growth only to 2.6% by 2075, and even "Moore's Law everywhere" makes income infinite only around 2060, because "it is only when the last weak links are automated away that the explosion fully unfolds" ([Jones and Tonetti](#)). Estimates for the coming decade sit near the slow end: at most 0.66% extra total factor productivity over ten years ([Acemoglu](#)), about half a point of growth a year ([Cowen](#)), and Anthropic's own 1.8 points a year falling to 0.6–0.8 once tasks are complements and success is counted ([Anthropic](#)). I read the view that inertia is the lesser problem as right about the destination and too relaxed about the decade, which is what matters to the workers who live through it.

15:46 Inertia barely moves the destination of the closing and decides who is paid while the gaps close: the owners of the slowest inputs, the trades that build them, and whoever writes the protocols on which others compete.

References

1. David 1990. <https://gwern.net/doc/economics/automation/1990-david.pdf>
2. Diego Comín, Bart Hobijn (2010). An Exploration of Technology Diffusion. *American Economic Review*. <https://www.aeaweb.org/articles?id=10.1257%2Faeer.100.5.2031>
3. nber.org, 2017. https://www.nber.org/system/files/working_papers/w24001/w24001.pdf
4. Erik Brynjolfsson, Daniel Rock, Chad Syverson (2020). The Productivity J-Curve: How Intangibles Complement General Purpose Technologies. *American Economic Journal Macroeconomics*. <https://www.aeaweb.org/articles?id=10.1257/mac.20180386>
5. St. Louis Fed. <https://www.stlouisfed.org/on-the-economy/2025/nov/state-generative-ai-adoption-2025>
6. BLS. <https://www.bls.gov/news.release/prod2.nr0.htm>
7. "harvest phase". <https://fortune.com/2026/02/15/ai-productivity-liftoff-doubling-2025-jobs-report-transition-harvest-phase-j-curve/>
8. MIT NANDA. https://mlq.ai/media/quarterly_decks/v0.1_State_of_AI_in_Business_2025_Report.pdf
9. Sultan, Farley and Lehmann 1990. <https://journals.sagepub.com/doi/10.1177/002224379002700107>
10. St. Louis Fed. <https://fredblog.stlouisfed.org/2026/08/does-generative-ai-save-time-at-work/>
11. NBER 2017. https://www.nber.org/system/files/working_papers/w23928/w23928.pdf
12. Michael R. Kremer (1993). The O-Ring Theory of Economic Development. *The Quarterly Journal of Economics*. <https://doi.org/10.2307/2118400>
13. "march of nines". <https://www.dwarkesh.com/p/andrej-karpathy>

14. Robin Dunbar (1992). Neocortex size as a constraint on group size in primates. *Journal of Human Evolution*.
[https://doi.org/10.1016/0047-2484\(92\)90081-J](https://doi.org/10.1016/0047-2484(92)90081-J)
15. Hannan and Freeman 1984. <http://www.iot.ntnu.no/innovation/norsi-pims-courses/harrison/Hannan%20&%20Freeman%20%281984%29.PDF>
16. Ege Erdil, Tamay Besiroglu (2023). Explosive growth from AI automation: A review of the arguments. *arXiv*.
<https://arxiv.org/html/2309.11690>
17. Quote Investigator. <https://quoteinvestigator.com/2019/01/03/estimate/>
18. Minneapolis Fed. <https://www.minneapolisfed.org/article/2026/ai-adoption-in-business-grows-steadily-but-unevenly>
19. IoT Analytics. <https://iot-analytics.com/state-of-enterprise-iot-from-iot-autonomous-connected-operations/>
20. ben-evans.com, 2026. <https://www.ben-evans.com/benedictevans/2026/9/3/ai-tools-and-transformation>
21. LBNL, Queued Up. <https://emp.lbl.gov/sites/default/files/2025-12/Queued%20Up%202025%20Edition%20-%2012.15.2025.pdf>
22. CEQ. https://trumpwhitehouse.archives.gov/wp-content/uploads/2020/01/20200612CEQ_EIS_Timelines_Report_Update.pdf
23. ABC. <https://www.abc.org/News-Media/News-Releases/abc-construction-industry-must-attract-349000-workers-in-2026-despite-macroeconomic-headwinds>
24. AGC. <https://www.agc.org/sites/default/files/users/user21902/2026%20Workforce%20Survey%20Analysis%20%284%29.pdf>
25. CSIS. <https://www.csis.org/analysis/genais-human-infrastructure-challenge-can-united-states-meet-skilled-trade-labor-demand>
26. Texas Tribune. <https://www.texastribune.org/2026/04/28/data-centers-texas-electricians-builders/>
27. Fortune. <https://fortune.com/article/nvidia-billionaire-ceo-jensen-huang-demand-for-gen-z-skilled-trade-workers-electricians-plumbers-carpenters-data-center-growth-six-figure-salaries/>
28. New York Times. <https://www.nytimes.com/2025/10/21/technology/inside-amazons-plans-to-replace-workers-with-robots.html>
29. Figure. <https://job-boards.greenhouse.io/figureai/jobs/4700899006>
30. Dexmate. <https://jobs.ashbyhq.com/dexmate/e92a4b08-1123-47f6-9d3d-3fc67fcd9df>
31. IDC. <https://news.cgtm.com/news/2026-01-24/IDC-report-China-leads-the-global-humanoid-robot-rise-in-2025-1KccOGZyVGM/index.html>
32. Omdia. <https://www.scmp.com/tech/tech-trends/article/3339346/chinese-firms-outpace-us-rivals-2025-humanoid-robot-shipments-agibot-takes-lead>
33. IFR. https://ifr.org/downloads/press_docs/Market_Presentation_WR_Press_Conference_2026.pdf
34. IDC. <https://www.idc.com/resource-center/blog/humanoid-robotics-commercialization-2026/>
35. Counterpoint. <https://counterpointresearch.com/insights/global-humanoid-robot-shipments-soar-nearly-300-percent-yoy-in-h1-2026>
36. IFR. <https://ifr.org/ifr-press-releases/news/five-million-robots-now-operate-in-factories-globally>
37. 24/7 Wall St. <https://247wallst.com/investing/2026/09/14/goldman-sachs-just-supercharged-its-humanoid-robot-prediction-5x-to-6-5-million-by-2035/>
38. heise. <https://www.heise.de/en/news/Optimus-Bot-Tesla-cancels-ambitious-production-targets-10742468.html>
39. Electrek. <https://electrek.co/2026/01/28/musk-admits-no-optimus-robots-are-doing-useful-work-at-tesla-after-claiming-otherwise/>
40. Tesla. <https://assets-ir.tesla.com/tesla-contents/IR/TSLA-Q2-2026-Update.pdf>
41. Electrek. <https://electrek.co/2026/09/25/tesla-optimus-production-ramp-hands-ai-generalization-problems/>
42. earnings call. <https://www.fool.com/earnings/call-transcripts/2026/08/05/tesla-tsla-q2-2026-earnings-call-transcript/>
43. SEC. https://www.sec.gov/Archives/edgar/data/1318605/000110465925108507/tm2530590d1_8k.htm
44. Figure. <https://www.figure.ai/news/production-at-bmw>
45. Boston Dynamics. <https://bostondynamics.com/news/boston-dynamics-opens-robotics-metaplant-application-center-to-train-humanoid-robots-for-manufacturing-tasks/>
46. " ~80% of Tesla's value will be Optimus". <https://www.cnbc.com/2025/09/02/musk-tesla-value-optimus-robot.html>
47. TechNode. <https://technode.com/2026/08/20/why-unitree-became-the-first-humanoid-robot-company-to-go-public-in-china/>
48. NDRC. https://en.ndrc.gov.cn/news/mediar/sources/202510/t20251021_1402139.html
49. Narayanan. <https://www.normaltech.ai/p/fact-checking-moravecs-paradox>

50. Valor. <https://valor.globo.com/financas/noticia/2025/04/29/caixa-e-bb-lideram-concentrao-do-sistema-financeiro-nacional-veja-nmeros.ghtml>
51. BCB. <https://www.bcb.gov.br/content/publicacoes/ref/202605/RELESTAB202605-refPub.pdf>
52. Resolução BCB nº 1. https://www.bcb.gov.br/content/estabilidadefinanceira/pix/Pix_Regulation/Resolution_BCB_1.pdf
53. BCB. https://www.bcb.gov.br/content/estabilidadefinanceira/pix/Regulamento_Pix/IX_ManualdeTemposdoPix.pdf
54. BCB. https://www.bcb.gov.br/content/estabilidadefinanceira/pix/relatorio_de_gestao_pix/relatorio_gestao_pix_2026.pdf
55. Febraban. <https://portal.febraban.org.br/noticia/3926/pt-br/>
56. Sarkisyan. <https://www.ssarkisyan.com/publication/instantpayments/>
57. Nubank. https://nu.com/media/2026/07/DataNubank-8_2026_Nubanks-presence-and-impact-across-Brazil.pdf
58. Nu Holdings. <https://www.businesswire.com/news/home/20260813187996/en/Nu-Holdings-Ltd.-Reports-Second-Quarter-2026-Financial-Results>
59. earnings call. <https://www.fool.com/earnings/call-transcripts/2026/08/20/nu-nu-q2-2026-earnings-call-transcript/>
60. BLS indexes, compiled by Chartive. <https://chartive.org/visualizations/what-got-more-expensive-cheaper-since-2000>
61. Anthropic. <https://www.anthropic.com/research/anthropic-economic-index-january-2026-report>
62. TrendForce. <https://www.trendforce.com/presscenter/news/20260930-13258.html>
63. TrendForce. <https://www.trendforce.com/presscenter/news/20260929-13255.html>
64. Nikkei. <https://asia.nikkei.com/business/technology/exclusive-tsmc-to-raise-chipmaking-prices-by-up-to-10-from-2027>
65. Altman. <https://blog.samaltman.com/three-observations>
66. working paper, September 2026. <https://basilhalperin.com/papers/singularities.pdf>
67. Jones & Tonetti 2026. https://web.stanford.edu/~chadj/JonesTonetti_Automation.pdf
68. Daron Acemoğlu (2024). The Simple Macroeconomics of AI. *National Bureau of Economic Research*. <https://www.nber.org/papers/w32487>
69. "the relevant constraint is the rate of drug approval, not the rate of drug discovery". <https://marginalrevolution.com/marginalrevolution/2025/02/why-i-think-ai-take-off-is-relatively-slow.html>