An Investment Thesis for the Economy Beyond Constraint

Arata Hoshino United Arab Emirates September 2026

Readership

This thesis is written for institutional investors, sovereign funds, venture and infrastructure investors, deep tech founders, and the people who design energy, industrial and population policy. It is a hypothesis. It analyzes, conceives and designs from the far side of the present constraint rather than from inside it, and it is meant to be used, as a hypothesis still, as a compass for investment and strategy.

Abstract

This thesis argues that contemporary civilization faces a single structural constraint, a stalled capacity to build, and that the constraint has three components: (i) the absence of effective energy abundance, (ii) a shortfall in labor input where automation has lagged, and (iii) the absence of institutions that draw in capital and talent. Economic stagnation and population collapse are symptoms of that constraint, not causes of anything.

The argument moves from the industrial revolutions of the past, through a diagnosis of the present, to the design of distribution and of value. The industrial revolutions are re-read along two axes, the efficiency with which matter is placed and the efficiency with which the future is predicted, and the asymmetry of the last thirty years becomes visible on those axes. The resulting supply constraint acts at two levels. At the macro level it narrows the range of large investments anyone conceives at all. At the level of the household it raises the prices of the goods automation never reached, and in doing so suppresses population: a second Malthusian trap. From that diagnosis the thesis proposes replacing the allocation of income by contribution to tasks with a distribution of ownership built on capital accounts opened at birth. In place of GDP, it draws on service-dominant logic, the theory in which value is created at use by the beneficiary rather than at exchange, to move the center of value to use-value, and sets out the conditions under which use-value could be measured.

On the far side of the constraint, where energy is abundant and automation accelerates under the right institutions, the capacity to conceive large things returns. People hold capital before they ever work and can give their time to creating value in the full sense rather than to productivity. The macro problems now feared, population collapse among them, recede as a by-product of the lifted supply constraint and the new mechanism of distribution, without treatment symptom by symptom, and supply and demand begin to lift each other.

The later sections place the world's clusters on the five terms into which the three variables decompose, and forecast the futures of two techno-blocs: an American one built on a division of labor and a Chinese one built on vertical integration. They argue for deepening the division of labor between the United States, the UAE and Northeast Asia, and for an ecosystem that integrates actors from all three. It then sets out what to invest in, term by term, states eleven predictions with the observation that would show each of them wrong, and closes with the institution to build.

A note on method

This thesis grew out of notes I kept in Japanese between 2018 and 2026. The ideas and the arguments are mine. I used Anthropic's Claude to help turn those notes into a single English text and to draw the figures from public data I had assembled. The errors are mine as well.

01 The capacity to build

Britain in 1847 poured almost 8 percent of that year's GDP into railways and more than quadrupled its network in a decade, on private judgment and private capital rather than any state plan.[1] The same country has now spent seventeen years and 46.8 billion pounds without opening a single mile of high-speed line.[2] What changed in the nearly two hundred years between was not national character or culture but the structure of abundance.

What a civilization can newly build in the physical world is the product of three terms:

Civilizational competitiveness C ≒ Institutions I × Effective energy abundance E × Labor input (automation) L

If any one of the three is near zero, the product is near zero, however high the other two. Competitiveness here means how fast a civilization can build the next physical substrate and how much new investment at scale it can carry. It says nothing about the quality of life of the people who live there, which is a different subject. Institutions do two jobs. They preserve returns, through property rights, investor protection, tax treatment and the speed of permitting, and that draws in capital. They permit failure and admit outsiders, and that draws in people. Effective abundance is a different thing from endowment, from having the resource under one's own ground. It means access to energy that is effectively free, and the belief that it will stay that cheap for a good part of the life of whatever is about to be built. Labor input, in the advanced economies of today, means automation, and it splits further:

L ≒ Deep tech T × Large capital K

Deep tech is whether the chokepoint technologies, the ones that stop everyone else the moment they are withheld, are held at home or within reach. Large capital is whether assets already accumulated can be moved at the scale that autonomy and computing substrate demand. Neither has to be built at home; both can be obtained through alliance or transfer, but only with somebody's permission.

The equation says nothing about the wealth a country already holds. It says how much the country can still build, and the constraint reaches further than what gets built. Options never reach the point of being weighed and rejected, because they stop being proposed. Where a premise is known to be fixed, no business plan is written on it, nothing reaches an investment committee, and nothing is conceived in the first place. The range of ventures we now feel to be realistic, and our common sense about the scale of investment, were both formed on the assumption that the three terms sit where they currently sit. So the constraint binds what can be conceived as well as what can be built.

Take two examples of societies that have held all three terms at once: Britain in the 1840s and the UAE today. In the UAE energy is cheap for three separate reasons: historically subsidized domestic gas, the 5.6 gigawatts of the Barakah nuclear plant, and Al Dhafra solar, which set a world record price in 2020.[3] Labor is carried by migrants who make up more than eighty percent of the resident population, and land, being desert, is effectively unlimited. The twenty-first-century UAE is pouring into computing substrate what nineteenth-century Britain poured into railways. Windows like these can also close, as Japan's did. Its postwar growth stood on cheap Middle Eastern crude; the window closed in 1973, on top of a slowdown already

[1] Andrew Odlyzko, "The collapse of the Railway Mania," rev. 2011: railway investment rose from 13 million pounds in 1845 to a peak of 44 million in 1847, "almost 8% of British GDP, and twice the military budget of that year," in sums actually spent at the prices of the day. Route length in Great Britain went from about 1,498 miles at the end of 1840 to 6,084 at the end of 1850 (B. R. Mitchell, *British Historical Statistics*, 1988, Table 5, p.541, which could not be obtained; the 1840 figure rests on a note in the 1962 predecessor and may be for the United Kingdom).

[2] Cumulative spending was 46.8 billion pounds at the end of March 2026, nominal, of which 44.2 billion is the current program and 2.6 billion the canceled second phase (National Audit Office, *High Speed Two reset*, HC 52, 2026; Department for Transport, *HS2 Six-Monthly Report to Parliament*, May 2026). The same report puts the cost at completion at 87.7 to 102.7 billion pounds and the opening of Old Oak Common to Birmingham between May 2036 and October 2039. HS2 Ltd began operating in January 2009.

[3] Barakah is four APR1400 units of 5.6 gigawatts; with all four in commercial operation from September 2024 it generated about 40 terawatt-hours in the following twelve months, up to 25 percent of national demand on an annual basis (Emirates Nuclear Energy Corporation). Al Dhafra solar, 2 gigawatts nameplate, was awarded in April 2020 at 1.35 cents per kilowatt-hour, the lowest price announced anywhere at the time, revised to 1.32 cents at financial close (Abu Dhabi National Energy Company, 22 December 2020). The IMF put implicit energy subsidies at 5.7 percent of GDP in 2011 (IMF Country Report No. 15/220).

under way in the domestic structure, and real growth halved from 9.1 percent a year to 4.2.[4]

The abundance of 1840s Britain rested in part on child labor and colonies, human costs that must not be repeated, and the labor cost of the present-day UAE rests on a global wage gap. The combination of energy abundance, fully automated manufacture and transport, and the right institutions would make the three conditions obtainable, for the first time in history, without human sacrifice, and anywhere. The abundance the nineteenth century bought with human pain, the twenty-first can buy back with machines. That is the civilizational meaning of committing capital to physical AI, meaning robots and automated plant, and to energy.

[4] Fiscal-year basis: an average of 9.1 percent from FY1956 to FY1973 and 4.2 percent from FY1974 to FY1990 (Cabinet Office of Japan, long-run national accounts). The Bank of Japan puts the slowdown principally on domestic structural factors, the fading of catch-up-driven productivity growth and the slowing of labor movement between sectors, with the oil shock laid over them (*Bank of Japan Review* 2012-J-17, December 2012; the October 2012 *Outlook* treats the same transition in Box 7, which could not be obtained).

02 Two convergence costs

Civilization advances in two directions on one substrate, energy, and the two directions compete for the same problems. For thirty years capital flowed to the direction that needed less energy, the digital one, and in the 2020s that direction ran into the wall of energy itself. The commitment of capital to physical AI and energy is therefore a return: capital going back to the side it had left behind.

The progress of civilization is made of one substrate, energy, and two directions that ride on top of it. Direction 1 drives to zero the cost of reaching the configuration of matter you were aiming at, an optimization measured in joules; transport, manufacture, construction, agriculture and the opening of space are all special cases of it. Direction 2 drives to zero the cost of reaching a predictive model of the accuracy you need, an optimization measured in bits; communication, computation, science, education and AI are all special cases of it. Both lower the cost of turning an intention into an outcome, one by changing the physical world and the other by knowing what to do.

To predict which means, which technology, a given problem will call for, treat the two directions as competing for it: the side with the lower convergence cost, the cost of turning an intention into an outcome, wins. To see which side that is, decompose the problem into steps and count the mass that must be moved at each. In shopping, search, comparison, ordering and payment all move zero mass; only the last leg, from warehouse to doorstep, moves any. So the problem mostly migrated to direction 2. Amazon beat the shopping robot people used to imagine on cost structure, not on the superiority of one vision over the other.

For thirty years capital and talent flowed almost one-sidedly into direction 2. It replicates at zero marginal cost, runs with low friction and needs little energy, and of those three properties only the last varies with the era; in a constrained world, capital flows toward the looser constraint. Then, in the 2020s, direction 2 ran into the wall of energy, and improved efficiency did not hold the wall off. A fall in convergence cost unlocks applications that were not previously viable, and demand grows by more than the efficiency gain. On the IEA's estimate, world data center electricity demand rose 17 percent year on year in 2025 and will roughly double by 2030, and AI-specific facilities alone rose 50 percent in 2025.[5] Over the same period the electricity consumed per AI task fell rapidly, and the saving went straight into more computation, used more widely, more often and for heavier purposes.

The movement of capital speaks more plainly than any forecast. In August 2025 the Palisades plant in Michigan was returned by the regulator from the decommissioning

[5] IEA, *Key Questions on Energy and AI*, 16 April 2026: data center electricity demand up 17 percent in 2025 and AI-specific facilities up 50 percent; 485 TWh in 2025 rising to about 950 TWh in 2030 on the central projection, for data centers as a whole. The IEA states that energy per AI task is falling by at least an order of magnitude a year but publishes no consistent annual series for it. The measured American series is Shehabi et al., *2024 United States Data Center Energy Usage Report*, LBNL-2001637: about 60 TWh a year through 2014 to 2016, 176 TWh in 2023, and 325 to 580 TWh in 2028 on the report's scenarios.

process to an operating license, the first time such a restart has been approved. "Operating" here is a regulatory category rather than the resumption of generation.[6] The decision is one sign of the structural change this thesis depends on: computing demand is making investment in energy supply politically unavoidable, and capital is already moving back.

03 A scarcity chosen

The slowing of growth across the advanced economies is a trap of its own kind, an advanced-country trap rather than the middle-income trap that stalls developing countries. Of the three terms, the one that broke was energy, and it broke by choice. In the 1970s the advanced countries stopped expanding their energy supply, and for the half-century since they have rebuilt their economies, institutions and values on the premise that energy does not grow. The source they turned to instead, one whose output does not follow demand, passed its cost to households through the electricity bill, and with it the number of children a household can afford.

The stagnation shows in three places. First, energy use per person stopped growing. In the United States it rose through the whole postwar period, peaked in 1978 and 1979, and has since fallen by close to a fifth; since efficiency gains have historically raised total consumption rather than lowered it, the fall reads as a halt to expansion, not as efficiency's reward.[7] Second, productivity growth broke in the same years. American total factor productivity grew at a steady pace from 1920 to 1970 and at about a third of that pace since, and a model that explains growth by the useful work extracted from energy fits American output until 1973 and 1974 and then starts to fall short.[8] Third, the country stopped building its energy supply: every one of the 41 reactors ordered in the United States after 1973 was canceled.[9]

No law of physics required that. American reactor costs were already climbing through the late 1960s and went far higher again for the reactors still under construction when Three Mile Island happened. Over the same years France held its escalation to a fraction of the American rate and Korea brought its costs down outright, both by building standardized designs over and over. The same technology escalated in one country and fell in another, and the difference was regulation and politics, and underneath both, whether a country built the same plant again or a new one every time.[10] The turn of the 1970s was chosen by the people of the advanced countries themselves, and nothing was prepared to replace the capacity given up. For the half-century since, the advanced countries have rebuilt their economies, their institutions and their values on the premise that energy does not grow. The scarcity described here is that rebuilding: a scarcity chosen, not a scarcity given.

The main new source the advanced countries turned to in that half-century was one whose output does not follow demand. The market value of wind falls from about 110 percent of the average power price near zero penetration to between 50 and 80 percent at a system share of 30 percent. The more wind there is, the more every other turbine is generating in the same hours, so it lowers the value of the electricity it produces. German wholesale power prices were negative for 573 hours in 2025, a record for the second year running.[11] A society that gave its baseload to a source suited only to supplement firm power passed the cost to households through the electricity bill. To choose a generation mix is to choose an electricity price curve; to choose an electricity price curve is to choose, as one bundle, the prices of housing, food and services; and to choose those prices is, in the end, to choose the number of children a household can have. If the whole of that chain had been printed on the

[6] The Nuclear Regulatory Commission's exemption from 10 CFR 50.82(a)(2), dated 24 July 2025 and effective 25 August 2025, allowed the permanently-shutdown certification to be withdrawn and fuel to be placed in the vessel again (NRC, ML25196A092). Holtec began loading fuel on 30 August 2026, and the NRC still showed zero output on 18 September 2026; no primary record of first criticality has appeared. The body relies only on the fact that a reactor was pulled back out of the decommissioning queue.

[7] U.S. Energy Information Administration, *Monthly Energy Review*, Table 1.7, series "Total Primary Energy Consumption per Capita": 207 million Btu per person in 1949, a joint peak of 350 in 1978 and 1979, and a preliminary 282 in 2025. The annual series begins in 1949, so nothing here reaches further back. The climb to the peak was not monotonic.

[8] Robert J. Gordon, *The Rise and Fall of American Growth*, Princeton University Press, 2016, figure 1-2 (TFP at 1.89 percent a year from 1920 to 1970 and about 0.65 percent from 1970 to 2014). Robert U. Ayres & Benjamin Warr, "Useful Work and Information as Drivers of Growth," INSEAD Working Paper 2002/121, in which the model reproduces American output to 1973 and 1974 and then slightly underestimates it; the authors attribute the shortfall to information technology capital rather than to energy. In their later book, *The Economic Growth Engine*, 2009, estimated over 1900 to 2005, the only significant break in the residual is at the Second World War.

[9] IAEA, *50 Years of Nuclear Energy*, GC(48)/INF/4, Attachment 3, 2004, para. 15, which takes the sentence from Cohn (1997) rather than from its own reactor records. The count is independently confirmed by the US Energy Information Administration, *Nuclear Plant Cancellations*, DOE/EIA-0392, 1983, Table 1, p.5: 41 units and 48,495 megawatts ordered after 1973. More than two thirds of all reactors ordered after January 1970 were canceled in the end.

ballot paper, it would have been a democratic choice; what was missing was any awareness that a choice was being made.

04 The household Malthus

This thesis holds that a second Malthusian trap is under way. In the world Malthus described, food production set the ceiling on population. In ours the ceiling is set by the household's economic capacity, and that capacity is governed by the equation in section 01. Where effective energy abundance and labor input stalled, automation never reached the goods a child needs, housing, education, childcare and healthcare, so their prices rose faster than everything else while the prices of automated goods collapsed. Households meet the constraint not as hunger but as a budget that does not stretch to a child, and falling fertility is the symptom of that constraint.

In the debate over falling fertility in the advanced economies, two explanations are standard: a change in values, and a failure of policy. People came to weight career and free time and to see children as a cost; or family policy is weak, and childcare, housing and education are inadequate. Both are partly right, and the question is whether either is a cause.

In the world Malthus analyzed, the constraining variable was food. Where land and agricultural technique have a ceiling, a rising population means less grain per head, and famine, epidemic, war, later marriage and contraception push the population back down toward that ceiling. That is the classical Malthusian trap. Nitrogen fixation by the Haber-Bosch process broke it in the early twentieth century: artificial fertilizer lifted yields so far that in the advanced countries a shortage of calories ceased, at the macro level at least, to bind. Malthus's schema was not wrong. What was wrong was the premise that agriculture would be the bottleneck forever. A population is still held down by the dominant physical constraint of its age. The constraint has simply moved, into the terms of the equation in section 01: effective energy abundance, labor input and institutions.

Nobody in an advanced country today gives "there is nothing to eat" as the reason for not having children, and fertility sits at historic lows all the same. The modern "food" is the household's economic capacity: whether, on its wages and at the prices it faces, a life with children is affordable at all. That capacity is itself a second-order variable, set by the energy structure and the level of technology. The price of a good or a service is set, roughly, by energy cost, manufacturing cost, transport cost, computing cost, and whatever labor cost has not been automated away. The last of these has a proper name in economics: Baumol's cost disease. When wages rise in a sector where productivity grows, the sectors hiring from the same labor market have to raise wages too or they cannot hire at all: education, healthcare, childcare, elder care. Their productivity has not risen and their wages have, so their prices keep rising.[12]

The cost of raising one child is made up almost entirely of Baumol sectors plus housing. Since 2000, American hospital services have risen more than 200 percent, college tuition more than 180 percent and childcare more than 100 percent, while the prices of televisions, toys and software have fallen by anything from tens of percent to more than 90 percent.[13] The goods automation reached collapsed in price; the person-to-person services it did not reach soared. Raising a child means buying, twenty years at a stretch, precisely the goods that automation has reached least.

[10] Jessica R. Lovering, Arthur Yip & Ted Nordhaus, "Historical construction costs of global nuclear power reactors," *Energy Policy* 91, 2016, 371–382. In overnight construction cost at 2010 prices, the 48 American units started between 1967 and 1972 went from roughly 600 to 900 dollars per kilowatt to roughly 1,800 to 2,500, and the units still under construction at Three Mile Island cost a median 2.8 times those finished before it; French costs rose from roughly 1,000 euros per kilowatt to 1,500 to 2,000 over 1971 to 1991; Korean costs fell by about half. The comparison is contested on its data range and cost definitions (Koomey, Hultman & Grubler, *Energy Policy* 102, 2017, and the authors' reply). Flyvbjerg's project database reads the same record as bespokeness rather than regulation, with nuclear plants overrunning by 120 percent on average everywhere (Bent Flyvbjerg & Dan Gardner, *How Big Things Get Done*, Currency, 2023, chapter 9 and Appendix A); the full version's Chapter 3 weighs the two readings.

[11] Lion Hirth, "The market value of variable renewables," *Energy Economics* 38, 2013, 218–236. The fall from 110 percent to 50–80 is a model result; the observed German values (1.02 at 2.0 percent penetration, 0.89 at 8.0 percent) are ex-post day-ahead outcomes, a different kind of evidence. Negative-price hours are the Bundesnetzagentur's SMARD figures: 457 in 2024, 573 in 2025.

[12] William J. Baumol, "Macroeconomics of Unbalanced Growth," *American Economic Review* 57(3), 1967, 415–426. The application to American education and healthcare is Eric Helland & Alex Tabarrok, *Why Are the Prices So Damn High?*, Mercatus Center, 2019, who make the mechanism central rather than sole. The sectors Baumol himself names are education, the police, hospitals, social services and inspection; childcare and elder care are the same mechanism applied later.

Figure 4-1 What got cheaper, and what did not

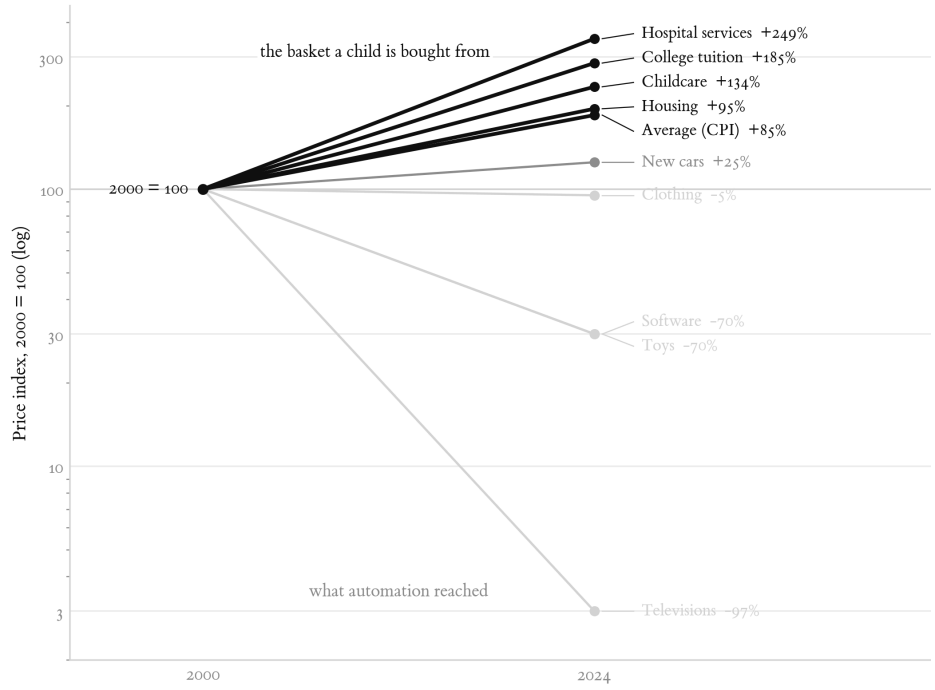

Log scale, so equal slopes are equal rates of change. Baumol sectors rise; automated goods collapse. Label positions on the right are spread for legibility; the line ends are exact.
Source: US Bureau of Labor Statistics, consumer price data, 2000–2024.

Figure 1. What became cheap and what did not (United States, 2000-2024). The scale is logarithmic, so equal slopes mean equal rates of change. The goods automation reached branch downward and the Baumol sectors branch upward. Everything a child needs is in the upper bundle.

The household's expected number of children is, in this reading, a function of one ratio:

Expected number of children $n^* \fallingdotseq g$(household disposable income $Y_d \div$ expected cost per child C_{child})

where C_{child} is a composite of the effective prices of housing, education, childcare and healthcare, each of them a Baumol sector, and each rising faster than anything else in the price system. That is the second Malthusian trap, the household Malthus of this section's title: a structure that suppresses population, derived mechanically from energy, technology and prices, and owing nothing to what anyone believes about children.

Falling fertility is therefore a price phenomenon, not a cultural one. Korea budgeted about 156 trillion won for low-fertility measures between 2006 and 2020, and 43 trillion more in 2021 alone, and the total fertility rate it arrived at in 2023 was 0.72, and 0.55 in Seoul.[14] Fifteen years of spending, on a budget line that grew more than fortyfold, left the price structure where it was. What separates fertility rates across the advanced countries now is not income but the cost of combining career and family.[15] Values are not a primary cause. They appear two or three steps down from a structure that sits below them: energy and technology set the price system, the price system sets the household's budget constraint and its anxiety about the future, and under that constraint households weight career over children as a rational principle of action, which later justifies itself as a value. To detect a change in a downstream

[13] Bureau of Labor Statistics CPI series, retrieved by the author through the API. From January 2000 to January 2024: hospital services +248.9 percent, college tuition +184.6 percent, day care and nursery school +134.3 percent; televisions about −97 percent, toys and computer software about −70 percent each. The electronics series carry hedonic adjustment, so the 97 percent is the price of a constant quality of television, not the cash price of the same object. That automation is what separates the two groups is this thesis's reading, not something the index series establish.

[14] National Assembly Budget Office, *Analysis and Evaluation of Programs Responding to Low Fertility*, August 2021: 155.63 trillion won of national funds budgeted over 2006 to 2020 across the three basic plans, with 42.90 trillion for 2021; annual budgets run from 1.03 trillion won in 2006 to 42.90 trillion in 2021. These are budgeted amounts, not settled expenditure, and the definition of a low-fertility measure widens over the period, so it is not one series on one definition. TFR is Statistics Korea, 2023 birth statistics (0.72 nationally, 0.55 in Seoul); the recovery to 0.75 in 2024 and a provisional 0.80 in 2025 is not used as evidence here, because a rebound cannot yet be separated from it.

variable and call it a primary cause is close to calling the movement of a thermometer needle the cause of the heat.

The place where this mechanism runs at the largest scale is China. Births in 2025 fell to 7.92 million, a drop of 17 percent in a single year and less than half the 2016 figure. That the decline will not reverse is already fixed by the number of people who have been born. The female population aged 20 to 34, which accounts for roughly 85 percent of births, goes from 105 million in 2025 to 58 million in 2050. Births spread from about age 18 to age 40, so every woman who will give birth in 2050 has been born by 2032 at the latest. The only degree of freedom left is the level of fertility, and holding it at its present level still takes births in 2032 down to around 6.5 million.[16] Of all the variables in this thesis, population is fixed earliest and moves last.

05 The abundance cascade

On a week away with 2GB of roaming data, we hold back as far as we can. Whether a thing is finite or infinite works on the mind out of all proportion to the actual scarcity. Latent business models branch like a tree, and to stop short because a constraint is in the way may be to cut away every branch that would have flowered further on. On energy, our civilization is that traveler with 2GB.

So a collapse in the cost of energy does more than make existing demand cheaper to serve; demand itself multiplies. Artificial light is one of the goods whose cost humanity has been driving down for longest, and on Nordhaus's measurement the real price of a thousand lumen-hours fell by 99.97 percent between 1800 and 1992.[17] When light became three thousand times cheaper, did humanity spend less on it? On the work of Tsao and colleagues, humanity has for three centuries, across six continents and five distinct lighting technologies, gone on spending about 0.72 percent of GDP on light. As efficiency rose, people consumed more light. They lit the city at night, ran factories around the clock, and later kept screens lit for most of the waking day. A share of income that stays in the same band while the price collapses is a rebound close to complete: people spend the whole saving from efficiency on more of the good. That is the reading taken here.[18]

Light is the clearest historical case, and the mechanism is not peculiar to lighting. Cheaper energy lowers the price of everything made with it; the lower prices unlock uses that were not viable before; the new demand pays for more supply, and energy gets cheaper again. That loop is the abundance cascade of this section's title, and what it unlocks next can be read off the cost structure. Between a quarter and two fifths of the cost of reverse osmosis desalination is energy, and energy accounts for 40 to 70 percent of the operating cost of vertical farming.[19] Launch cost has fallen twentyfold, from roughly 54,500 dollars per kilogram on the Space Shuttle to roughly 2,720 on Falcon 9, and Starship is spoken of as aiming two orders of magnitude below that again, a target of ten dollars a pound, about 22 dollars a kilogram, that Jones attributes to Musk.[20] Complete the third order of magnitude and earth orbit becomes an industrial location.

Supply arrives on a manufacturing timetable, and prices move in the 2040s. Ranked by how much firm power they can add in the long run, the sources are fusion at the top, then conventional large reactors and small modular reactors (SMRs) built in series in a factory, fossil fuels as a bridge, and renewables as a supplement on top of firm power. What brings the price down is how many units get built. The cost curve follows Wright's law, under which unit cost falls at a constant rate for every doubling

[15] Gary S. Becker, "An Economic Analysis of Fertility," 1960. Matthias Doepke, Anne Hannusch, Fabian Kindermann & Michèle Tertilt, "The Economics of Fertility: A New Era," NBER WP 29948, 2022. Price explains the conditions under which fertility falls, not those under which it rises: the Nordic decline since the 2010s under the most generous family policy in the world (Hellstrand et al., *Demography* 58(4), 2021) and Israel's secular fertility near 2 in a high-price society (Central Bureau of Statistics, *Statistical Abstract of Israel 2025*, Table 2.28: 2.87 overall, 1.98 among the non-religious) mark the edge of this explanation, and no record exists of a state in which the constraint has come off, so the prediction in the body is conditional rather than a law.

[16] National Bureau of Statistics of China, 2025 statistical communiqué, 28 February 2026 (7.92 million births in 2025, 9.54 million in 2024, a fourth consecutive annual fall in population). The projection to 2050 is this thesis's own cohort component method, applying an age-specific fertility distribution to cohorts already born and calibrated on the 2025 actual. The National Bureau of Statistics does not publish a total fertility rate; the 0.93 used in the projection is a private estimate. The population base is robust, since every woman who will give birth in 2050 has been born by 2032. The absolute number of births is not.

[17] William D. Nordhaus, "Do Real-Output and Real-Wage Measures Capture Reality?" in *The Economics of New Goods*, University of Chicago Press, 1996, ch.1. The real price of a thousand lumen-hours fell from 429.628 cents in 1800 to 0.124 in 1992, both in 1992 dollars (Table 1.4, p.46). This is the series for the best available technology, not for the average household.

[18] J. Y. Tsao et al., "Solid-state lighting: an energy-economics perspective," *Journal of Physics D* 43(35), 2010, 354001, on the data of Tsao & Waide, *LEUKOS* 6(4), 2010, 259–281: three centuries, six continents, five lighting technologies; 0.39 to 1.30 percent of GDP across seventeen observations, best fit 0.72. Neither paper puts a number on the rebound; the step from a stable expenditure share to a rebound close to complete is this thesis's own, though Fouquet & Pearson (*EEEP* 1(1), 2012) find backfire in British lighting directly, with energy used for light rising after nineteenth-century

Figure 5-1 Three centuries of light: the purest case of the cascade

The three points on the right are the empirical settings the paper reports, three centuries apart, from tallow candles to grid electricity. The shaded band is the paper's range of 0.39 to 1.30 percent of GDP across seventeen observations; the line through it is its best fit of 0.72, not a value measured three times over.
Source: Left: Nordhaus (1996), in *The Economics of New Goods*. Right: Tsao et al. (2010), *J. Phys. D* 43(35), 354001.

Figure 2. Three centuries of light. The price fell by a factor of three thousand, and spending on light stayed in the same band of income throughout, 0.39 to 1.30 percent of GDP across seventeen observations with a best fit of 0.72. The three points are the settings the paper reports, the United Kingdom in 1700, the off-grid world in 1999 and the on-grid world in 2005, spanning tallow candles to grid electricity.

of cumulative production: solar module prices fell 20 to 25 percent per doubling, and by more than 80 percent in the 2010s alone.[21] SMRs aim to convert a reactor from a civil engineering project built on site into a product manufactured in series, which would put nuclear on a manufacturing learning curve rather than a construction one. But between the ignition of the first fusion machine and a tenfold reduction in cost per unit lie several decades and several hundred reactors, which is why the effect on price arrives in the 2040s.

One caution about the supply side. How fast energy actually becomes abundant is decided not at the generator but at the connection between generator and load: the queue for new grid connections, the factories that make transformers and switchgear, the land and permits for transmission lines and substations. Generation has grown exponentially;[22] the grid connection has not. In the United States the median project entering operation in 2025 had waited five years from application and only 13 percent of applications ever reach operation,[23] and a large transformer takes two to three years to deliver, two to three times what it took in 2010.[24] However many factories are added, each of these lead times stays measured in years. Two things follow for an investor. Prices in the transition rise before they fall, because demand explodes while supply is held back. And the scarce asset is not the generating plant but the site already connected to the grid, and the position in the queue. Electricity does not travel far and land does not travel at all, so effective energy abundance is the one term of the equation that cannot be shared among partners, and the choice of site is itself the investment decision.

The cascade also runs back through the household, and this is where section 04 and this section meet. The goods that priced children out were the goods automation had not reached, and the reason it had not reached them was the same shortfall in energy and labor input. There has only ever been one cure for Baumol's cost disease: raise productivity in the sick sector directly, with technology. When cheap power and physical AI reach housing, education, childcare and healthcare, the one part of the

efficiency gains.

[19] Desalination: Voutchkov (*Desalination* 431, 2018) puts energy at 25 to 40 percent of a cost of about 1.1 dollars per cubic meter, and Ghaffour, Missimer & Amy (*Desalination* 309, 2013) at 19 to 40 percent for seawater reverse osmosis. Vertical farming: Pennisi and colleagues (*Agronomy for Sustainable Development* 45, 2025) give lighting at 42 to 80 percent of energy use and climate control at 16 to 43 percent. The share of operating cost is where published figures pull apart, from about seven tenths in a paper design on German tariffs with capital excluded (Banerjee & Adenaeuer, 2014) to 12 to 25 percent in a survey of operating North American farms (Agrilyst, 2017); the gap is mostly labor cost and the price of electricity.

[20] Harry W. Jones, "The Recent Large Reduction in Space Launch Cost," ICES-2018-81, 2018. Measured instead by delivery to the ISS, the Shuttle's 105,800 dollars per kilogram against roughly 25,000 for Falcon 9 and Dragon makes the gap four times rather than twenty.

[21] T. P. Wright, "Factors Affecting the Cost of Airplanes," *Journal of the Aeronautical Sciences* 3(4), 1936, 122–128. The solar series is Our World in Data (CC BY 4.0, after Nemet 2009 and IRENA); on it the module price falls 81.4 percent between 2010 and 2019 and 85.3 percent between 2010 and 2020, and the text takes the more conservative of the two. That a rapid transition is cheaper is Way, Ives, Mealy & Farmer, *Joule* 6(9), 2022.

[22] Cumulative installed capacity of 0.54 megawatts in 1975 and 1,866,898 megawatts in 2024 is the Our World in Data series (IRENA basis, grid-connected) as it stood in September 2026; an earlier vintage put 2024 at 1,852,359 megawatts. China is the National Energy Administration's figure of 1.84 billion kilowatts of wind and solar at the end of 2025. The 47.3 percent share is computed here against the total installed capacity reported for the same date, and is not published by the administration.

[23] Lawrence Berkeley National Laboratory, *Queued Up: 2026 Edition*: for projects entering operation in 2025 the median period from application to operation was 61 months; of applications filed between 2000 and 2020, 13 percent by capacity reached operation and 75 percent were

price system that only ever rose begins to fall, the ratio in section 04 turns, and the second Malthusian trap loosens. Population can then turn upward, and with it the size of the market as a whole. Abundant supply does not only multiply demand for what already exists. It restores the number of people who will demand anything at all. How long that restoration takes is the subject of the next section.

06 The valley of demand

The far side of the constraint is a larger economy: abundant supply, a recovering birth rate, and a demand market that grows again. The problem is the gap in time between supply and demand. Once the constraint comes off, supply grows on a timetable of years. Demand grows on a timetable of generations, and a birth rate that turns upward in 2035 adds consumers at once but adds workers, taxpayers and founders only decades later. This thesis calls that gap the valley of demand, and it lasts about half a century. Through it, the returns that can be relied on are the ones that do not scale with head-count: stakes in automation capital itself, and the infrastructure of energy and computation.

The gap cannot be closed faster because an economy needs at least four head-counts, labor supply, demand, risk-takers and taxpayers, and machines replace only the first. If nobody wants the thing, then however much you make, it is inventory. How many people will put their own lives into a business nobody has tried depends on the size of the population, not on productivity. You can tax the wealth machines generate, but a human community is what carries that taxation and maintains its institutions. And the four come back at different speeds. A person is a consumer from birth. On the age profiles of the National Transfer Accounts, which tabulate income and consumption by age, their labor income does not exceed their consumption until about 30.[25] The founders of the fastest-growing firms are on average in their mid-forties.[26] Even if the constraint came off completely in 2035 and fertility in the high-income countries jumped to replacement that same year, which cannot happen, the valley would still last into the 2080s.[27]

Three moves bridge the valley. The first is the portfolio rule above: hold returns that do not depend on head-count, and treat businesses whose returns are proportional to the volume of demand as facing a structural headwind for the whole period. The second is to build demand where the people are. In those decades population is still rising mainly in Africa, the Levant, Pakistan, and parts of Central Asia and Central America, where purchasing power per head is low. Demand is population multiplied by purchasing power, so those regions hold future demand rather than present demand, and for the advanced economies holding plant they cannot sell, building up the ability of those populations to buy is less charity than the work of turning that plant from a liability back into an asset. The third move is the largest and has a section of its own: to make sure that when demand does return, it returns to households that own a share of the capital producing for them.

withdrawn. These are generation-side statistics; the load-side queue is not covered.

[24] Wood Mackenzie, *Making the connection*, Executive Summary, September 2025, slide 5: generator step-up transformers at 143 weeks after receipt of order and power transformers at 128 in the second quarter of 2025, for the American market. The baseline is the Department of Energy's *Large Power Transformers and the U.S. Electric Grid* (April 2014): five to twelve months from domestic producers and six to sixteen from abroad in 2010, so the 2025 figures are a doubling to a tripling.

[25] National Transfer Accounts Project, per-capita profiles by age; the 2011 United States profile is used here. On that profile labor income first exceeds consumption at 30, and labor income per head at 20 to 24 is 30 percent of its peak at 49, reaching half of peak at 26 and 80 percent at 33. This is a cross-section of a single year rather than a cohort followed through time.

[26] Pierre Azoulay, Benjamin F. Jones, J. Daniel Kim & Javier Miranda, "Age and High-Growth Entrepreneurship," *American Economic Review: Insights* 2(1), 2020, 65–82. Administrative data on 2.7 million founders of American employer firms: 41.9 years on average across all founders, 45.0 for firms in the top 0.1 percent by growth.

[27] The total fertility rate is stepped from 1.42 to 2.1 in 2035 and held there. The base population is World Bank data for high-income countries, 2025, run through a cohort component projection with a skew-normal fertility schedule (mean age 29.0), a Gompertz-Makeham survival schedule held fixed, and migration set to zero. It is a model rather than a forecast: zero migration makes the valley deeper than it would be, and what survives is the timing, since the year of the floor stays within 2085 to 2090 when life expectancy is moved four years either way.

Figure 5-5 The valley of demand, on the most optimistic assumption available

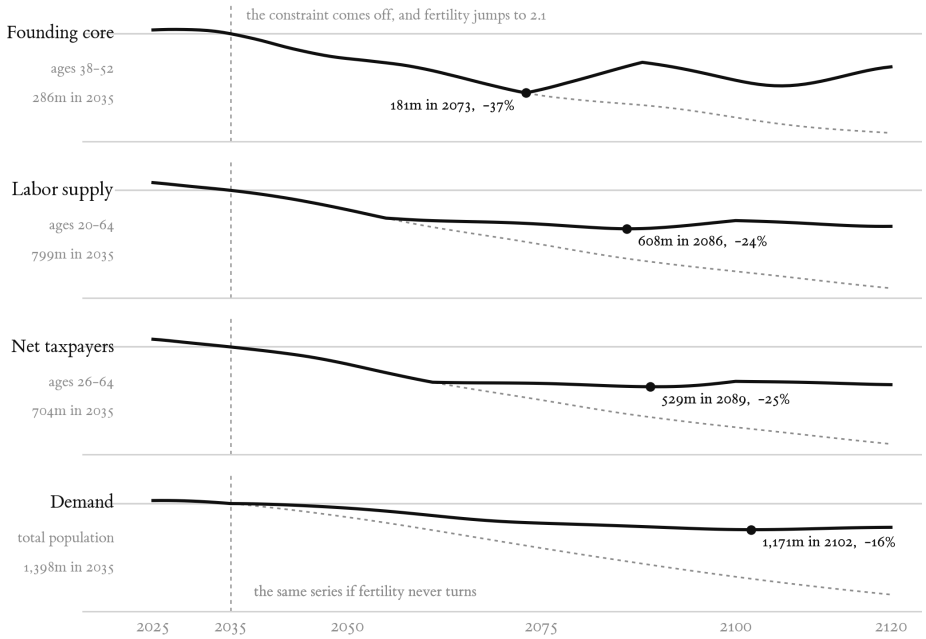

High-income countries, each row indexed to its own 2035 level. The total fertility rate is held at 1.42 and made to jump to 2.1 in 2035, with zero migration; no faster reversal is possible, so the valley drawn here is a lower bound. The dotted line is the same series without the reversal. The four floors do not fall in the same year.
Source: Author's cohort component projection; base population from World Bank data by age and sex (2025).

Figure 3. The valley of demand, on the most optimistic assumption available. Each row is indexed to its own 2035 level. The four floors do not come in the same year, and the founding cohort, the age group that starts companies, falls furthest.

07 The distribution of ownership

Modern economies distribute wealth through tasks. Work is broken into tasks, and income is paid for performing them. That mechanism is breaking, because automation is taking the tasks, and the capital that automation creates, and that abundant energy will create after it, has to be distributed some other way. This thesis proposes an expanded form of ownership distribution: every citizen holds a stake in automated capital from birth, before they ever work. In one line, every citizen becomes an investor before becoming a worker. The earliest examples already exist, in the accounts the United States opens for every child born from 2025 to 2028 and in Connecticut's baby bonds. The proposal is to take that design to the scale the coming decades require, and to say where the capital comes from.

Tasks are failing as the instrument. The labor share has fallen since the early 1980s across most countries and industries, and about half of the fall is explained by machines and computers getting cheaper and firms replacing people with them. Automation always has two effects, displacement, in which a machine takes a human task, and reinstatement, in which new tasks are created for people, and in the United States of recent years displacement accelerated while reinstatement weakened.[28] The break is already observable. In the occupations AI reaches most, employment of young workers has fallen relative to other age groups while employment of experienced workers has held, through a halt in hiring rather than layoffs, and the gap has widened with each release of the data.[29] The link between task and income is breaking at the entrance, at the young end.

[28] Loukas Karabarbounis & Brent Neiman, "The Global Decline of the Labor Share," *Quarterly Journal of Economics* 129(1), 2014, 61–103. Daron Acemoglu & Pascual Restrepo, "Tasks, Automation, and the Rise in U.S. Wage Inequality," *Econometrica* 90(5), 2022, 1973–2016. The "about half" decomposition is contested on its mechanism rather than its size: Elsby, Hobijn & Şahin (*BPEA*, Fall 2013) and Bergholt, Furlanetto & Maffei-Faccioli (*AEJ: Macroeconomics* 14(3), 2022) find automation rather than cheaper machines doing the work.

The familiar remedies do not restore it. The political demand for redistribution is rising,[30] and the instruments it reaches for first have a record. Around 1990, twelve OECD countries taxed net wealth; three still do. The reason is not that high rates blunt the will to work but that the tax base moves: where the base and its holders can move, taxing wealth country by country has not been sustained.[31] Cash transfers help the recipient without changing the structure. In the largest experiment to date, people given a thousand dollars a month for three years worked slightly less, were slightly less likely to be employed, and reported gains in stress and food security that faded after the first year.[32] Cash buys autonomy and immediate peace of mind. A cheque does not change who owns the capital, and ownership is the variable that is moving.

The stake works at every scale already tried. The American accounts pay 1,000 dollars into an index fund for every child born between 2025 and 2028, and Connecticut's baby bonds pay 3,200 dollars for every child born on Medicaid. Alaska has paid every resident a dividend from oil for four decades with no measurable effect on employment, and Norway's sovereign fund is collective capital of roughly 380,000 dollars per citizen.[33] Broad middle classes have repeatedly been made by turning citizens into owners: the Homestead Act moved about a tenth of American territory into private hands, and postwar Japanese land reform moved 1.74 million hectares to 4.75 million tenant households.[34] The nineteenth century distributed land, the twentieth distributed home ownership and pensions, and the twenty-first has to distribute a stake in automated capital, because as the labor share falls, households can hold their share of income only by owning capital.

Who pays in, and from what. The answer has to be consistent with the diagnosis: a stake funded mainly by taxing a mobile base would rebuild the failure it is meant to escape. Three sources are consistent with it. The first is public seed capital of the kind already in force, small, universal and paid at birth, so that compounding does the work over the twenty years before the stake is needed. The second, and the largest, is a share in the new capital itself. The scarce inputs of the coming build-out are the ones the state controls, the connected site, the permit, the land and the water. A state that grants them can take equity in the energy, computing and automation plant it enables, as Norway did with its oil, and place that equity in a citizens' fund rather than in the treasury. The third is voluntary: contributions from firms and wealthy individuals, whose interest is not only philanthropic, since the same firms will hold the plant that needs a market in the valley of the previous section, and capitalizing the next generation of consumers is the cheapest way to create one.

The same technology also changes who can own a business. Owning a business used to mean hiring people and taking responsibility for their lives, and that fixed cost set the minimum scale at which a firm is viable. AI and robotics lower that minimum scale. But if the threshold falls without any design for dispersing ownership, what multiplies is people who run a business but own no capital, and the stake at birth is the initial capital that lets such a person become an owner.

The reason to hurry is that institutions and population run on different clocks. Institutions of distribution can be made by legislation, so once the decision is taken they move within a few years. Population cannot be legislated into place: a child born now is a consumer from the first day, takes twenty years to reach the tax rolls and thirty to earn more than they consume. To institutionalize distribution in the 2030s is to buy the demand, the challengers and the tax revenue of the 2050s now.

[29] Erik Brynjolfsson, Bharat Chandar & Ruyu Chen, "Canaries in the Coal Mine?" Stanford Digital Economy Lab, 2025, revised 2026. The authors note this is correlation rather than causation, and the figure has moved with the data: 13 percent in August 2025, 16 in November 2025, 19 in the June 2026 follow-up. In Danish registers, Humlum & Vestergaard (2026) bound the part attributable to AI at 0.14 of a decline of about 0.6 percentage points, while Hosseini & Lichtinger (2026) find junior employment down 9 percent at adopting firms.

[30] On 4 November 2025, Zohran Mamdani won the New York mayoralty and Katie Wilson the Seattle mayoralty. The polling is Jeffrey M. Jones, "Image of Capitalism Slips to 54% in U.S.," Gallup, 8 September 2025 (54 percent favorable toward capitalism among all adults, 39 percent toward socialism).

[31] OECD, *The Role and Design of Net Wealth Taxes in the OECD*, 2018. Sarah Perret, *Fiscal Studies* 42(3–4), 2021, 539–563. The record is no proof that taxing wealth is wrong, only that it has been hard to sustain under mobility. In Scandinavian data a one-point rise in the top wealth tax rate reduces the stock of wealthy taxpayers by about 2 percent (Jakobsen, Kleven, Kolsrud, Landais & Muñoz, NBER WP 32153, 2026), while in Switzerland a one-point cut raises reported wealth by at least 43 percent, only a quarter of it through people moving (Brülhart, Gruber, Krapf & Schmidheiny, *AEJ: Economic Policy* 14(4), 2022). The base moves through exemptions, valuations and evasion before it moves house.

[32] Eva Vivalt, Elizabeth Rhodes, Alexander W. Bartik, David E. Broockman, Patrick Krause & Sarah Miller, "The Employment Effects of a Guaranteed Income," NBER Working Paper 32719, 2024, revised 2026: hours worked fell by 1.3 a week and employment by 4.2 points (standard error 1.7) against a control group of 2,000 on 50 dollars a month. On health, Miller and colleagues, NBER Working Paper 32711, 2025. Both concern transfers of limited duration, which do not mean the same thing to a recipient as capital income expected to last a lifetime.

[33] One Big Beautiful Bill Act (Public Law 119-21) §70204: 1,000 dollars for each United States citizen child born between 1 January 2025 and 31 December 2028, invested in funds

08 Moving the center of value

The last element of the design is value itself. It is the most abstract part of this thesis, so it is stated in four steps. First, most of the value people create never enters GDP. Second, that is no accident: GDP counts prices, and price was adopted two and a half centuries ago as a stand-in for a value that could not then be observed. Third, the two jobs price does, rationing scarce things and standing in for value, are both being taken over, one by cheap energy and the other by cheap measurement. Fourth, as that happens, what an investor buys changes, from a share in the company that makes a thing to a claim on the benefit the thing delivers to the people who use it.

Step one: what value is, and how much of it GDP misses. To create value is to apply one's own competence for the benefit of another.[35] Under that definition, serving coffee in a cafe and a mother caring for her child are the same act, and the only difference is that one carried a price. How much GDP misses can be measured, and three different measurements agree. Unpaid household production is worth about a quarter of GDP in the official satellite accounts. Paid work and study fill under five hours of the average day, while unpaid work and leisure fill more than eight.[36] And the median internet user would need roughly fifty dollars a month to give up one free platform, a good with a near-zero contribution to GDP.[37] The value that market transactions count is a minority of the value created in a day.

Step two: why GDP is built that way. Adam Smith drew the line between the two. In Book I, Chapter IV of *The Wealth of Nations* he distinguished use-value, what a thing is worth to the person who uses it, from exchange-value, what it fetches in trade. He noted that nothing is more useful than water and almost nothing can be bought with it, and then confined everything that followed to exchange-value, because exchange-value could be measured and use-value could not. He called use-value the real value. The paradox was never solved; it was set aside, and economics later forgot that it had been.[38] GDP is that decision made into an institution: it counts what was paid for, not what was of use.

Step three: why the stand-in is weakening. Price does two jobs: it rations scarce things, and it stands in for value. Cheap energy takes over the first. In information goods, where one more copy costs almost nothing, price has already stopped signaling scarcity, and the collapse in the cost of energy and automation extends the same mechanics into the material world. Measurement takes over the second, more slowly. Price stood in for value because in Smith's time there was no way to observe whether a benefit had been felt. That is no longer true. Behavioral and biometric data are already recorded, and in principle the response of the nervous system itself can be. If the fact that a beneficiary felt a benefit becomes recordable, the reason for keeping a stand-in disappears. Of the two jobs the second is the harder to displace, because it needs not only measurement but trust that the measurement was not faked.

Step four: what changes for money, and for investors. Money loses its first job, rationing, before it loses its second, measuring value, and the units that take over rationing are units of energy. The token that counts AI inference and the carbon emission allowance were both designed as units of account, and both are converging on currency; anything issued, exchanged and stored as a unit of a scarce resource is doing part of money's work. For an investor the change is this. Today one funds a company and is paid out of its profits, which is to say out of the exchange-value it captures. Once the benefit a thing delivers can be measured, one can fund the thing and be paid according to the measured benefit it delivers to its users, whoever made

tracking a broad index of American equities; the accounts went live on 4 July 2026. Connecticut's Baby Bonds: 3,200 dollars for each child born on Medicaid on or after 1 July 2023, usable between 18 and 30 for a home in the state, a business, higher education or retirement saving. Alaska: dividend paid annually since 1982; Damon Jones & Ioana Marinescu, *AEJ: Economic Policy* 14(2), 2022, 315–340. Norway: NBIM, *Annual report 2025*.

[34] On the Homestead Act, the US National Archives and National Park Service (more than 1.6 million applications processed through 1934, moving more than 270 million acres into private hands). Japanese land reform is the Ministry of Agriculture's report to ICARRD (1.74 million hectares, 4.75 million tenant households); the fall in the tenanted share, from 45.9 percent to 9.9, is Chiang, Fan & Hsu, IZA DP No. 18095, 2025.

[35] Stephen L. Vargo & Robert F. Lusch, "Evolving to a New Dominant Logic for Marketing," *Journal of Marketing* 68(1), 2004, 1–17; the axiomatization is "Institutions and axioms," *JAMS* 44(1), 2016, 5–23; the book is *Service-Dominant Logic*, Cambridge University Press, 2014.

[36] Benjamin Bridgman, Andrew Craig & Danit Kanal, "Accounting for Household Production in the National Accounts," *Survey of Current Business* 102(2), 2022 (22 to 25 percent of GDP). Time use is the OECD Time Use Database and Veerle Miranda, OECD Working Papers No. 116, 2011, averaged over 29 countries and ages 15 to 64: 4.6 hours a day of paid work and study, 3.4 to 3.5 of unpaid work, about 4.8 of leisure, over all days including weekends and both employed and non-employed, which is why paid work looks short.

[37] Erik Brynjolfsson, Avinash Collis & Felix Eggers, "Using massive online choice experiments to measure changes in well-being," *PNAS* 116(15), 2019, 7250–7255. The median monthly compensation demanded for giving up one free platform was about 48 dollars.

[38] Adam Smith, *The Wealth of Nations*, 1776, Book I, Ch. IV; the definition of real value is at pp.30–31. Aristotle separated use-value from exchange-value in the fourth century BCE, and the medieval Schoolmen held that use-value came first. The reading

it. The claim attaches to the use of what was made, not to the organization that made it. Section 10 lists the infrastructure that would make such a claim enforceable. Without it, neither this section nor the distribution of the previous one has a unit to settle in.

follows Lusch & Vargo, 2014, ch.2.

09 The center of gravity has gone to sea

Three conclusions, and then the reasoning. The unit of the world economy is the cluster, not the nation. The clusters are concentrated along the rim of the Indo-Pacific, because the heavy goods of the physical build-out move by sea, and this thesis therefore works with nine clusters, six of them on or near that rim and three, Texas, BosWash and Loxbridge, tied to it by capital and trade: the San Francisco Bay Area, Texas, BosWash (the Boston-Washington corridor), the UAE, Northeast Asia (Japan-Korea-Taiwan), Loxbridge (the London-Oxford-Cambridge triangle), the Yangtze Delta, the Pearl River Delta and Bengaluru-Hyderabad. And when the nine are measured against the five terms of the equation, no place carries every term, so a cluster's standing is decided by how it reaches the terms it lacks: a division of labor among clusters that co-create beats vertical integration inside one country.

The unit and the sea first. Divide the world's land by the continuity of its lights at night and the economy's concentrations appear by themselves. The 40 megaregions extracted this way hold under a fifth of the world's population and produce two thirds of its output and most of its patents.[39] Agglomeration has not come apart as the cost of communication approached zero, because what density buys, in labor pools and in learning from the people nearby, does not travel down a wire. And the clusters are where the sea is. More than eighty percent of world merchandise trade by volume moves by sea,[40] and the physical substrate treated here, transformers, turbines, reactors and building materials, is all on the heavy side, so distance from a port is distance to construction. The cities along the rim of the largest ocean on earth grew large through geography rather than effort. Land within 300 kilometers of the Indo-Pacific is 12 percent of the world's land, and by 2050 it carries a third of the world's people and close to two fifths of its output.[41] The littoral states of that ocean taken together already produce two thirds of world GDP, up from about half in 1980.[42]

The analysis: only two techno-blocs assemble all five terms, one by division of labor and one by vertical integration, and the second is bottlenecked by institutions. Seen at the level of the cluster, the equation splits into five terms: institutions, effective energy abundance, deep tech in bits (software and AI), deep tech in the physical (machines, materials and semiconductors), and large capital. A cluster's standing is set by whether it can reach the terms it lacks, not by how many it holds. This thesis calls a set of clusters that reach one another's terms a techno-bloc, and at present only two have all five assembled. The American techno-bloc assembles them by division of labor. The Bay Area specializes in bits, Japan-Korea-Taiwan in the physical, and Texas, BosWash and the UAE in institutions, energy and capital, and laid over one another they very nearly cover all five terms. The Chinese techno-bloc tries to assemble them by vertical integration inside one country. Its advantage is speed of decision and resilience against a cut in supply. Its weakness is that any one term that is structurally low becomes the bottleneck for the whole, and in the Chinese case that term is institutions. Direct investment from abroad has turned down amid economic-security concerns,[43] and for outsiders the country is hard to enter and hard to stay in: a foreign professional's right of residence is tied to the employer, and language, culture

[39] Richard Florida, Tim Gulden & Charlotta Mellander, "The rise of the mega-region," *Cambridge Journal of Regions, Economy and Society* 1(3), 2008, 459–476: 40 megaregions with 17.7 percent of world population, 66 percent of output and 85.6 percent of patents. Output is estimated from the intensity of night-time lights and renormalized to national GDP, a spatial reallocation of world GDP rather than an independent measurement. The data are from around 2000.

[40] UNCTAD, *Review of Maritime Transport 2017*, p. x: more than 80 percent of world merchandise trade by volume and more than 70 percent by value is carried by sea. The 2024 and 2025 editions do not state this ratio, so the 2017 edition is cited.

[41] This thesis's own aggregation on a two-minute mesh. Islands are counted with the state that holds them; Australia, New Zealand, Papua New Guinea and Madagascar are excluded. The coastline is SRTM15+ v2.7, borders are Natural Earth 10m, the within-country breakdown is Kummu, Taka & Guillaume, *Scientific Data* 5, 180004, 2018, and the projection is Murakami & Yamagata, *Sustainability* 11(7), 2106, 2019 (SSP2). Holding the ratios at their 2015 values gives 34 and 38 percent.

[42] This thesis's own aggregation from the IMF *World Economic Outlook* database, October 2025: the littoral states of the Indo-Pacific hold 69 percent of world population and produce 66 percent of world GDP at purchasing power parity, against 53 percent in 1980. At market exchange rates the figure is 65.5 percent. The 1980 value rests on the 49 of the 67 countries and territories for which data exist. The largest single term is the United States; excluding it takes the figure down to 51.7 percent.

Figure 8-1 The Indo-Pacific coast, where most of the nine clusters sit

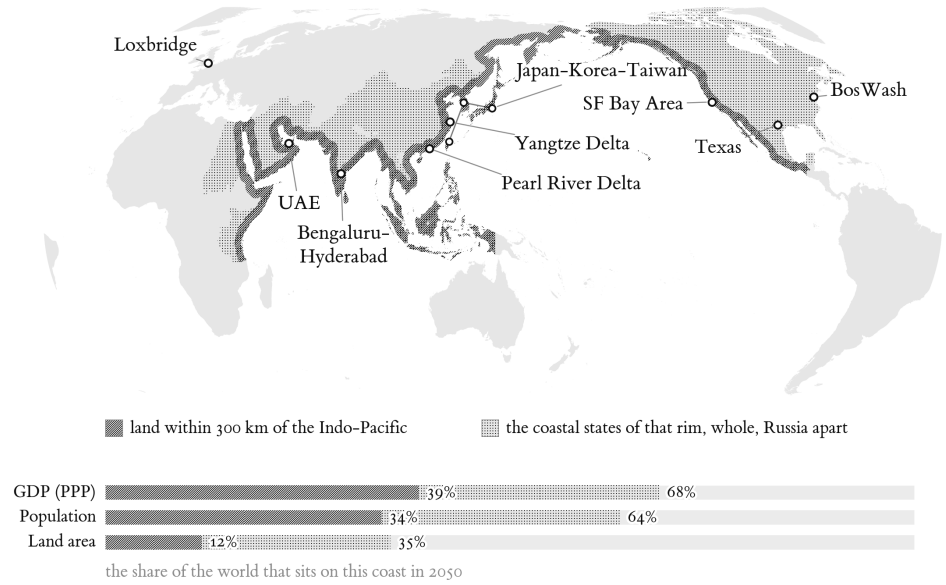

Two readings of the same rim. The black diagonals are land within 300 km of the Indo-Pacific, the Pacific and Indian oceans with the Gulf and the Red Sea, from the Bering Strait down to Mexico's southern border on one side and Tanzania's on the other. Islands go with the state that holds them; Australia, New Zealand, Papua New Guinea and Madagascar are out. The mesh is the coastal states of that rim taken whole, Russia apart because its body lies inland. The three clusters outside the band are coastal too, facing the Atlantic. Each bar shows the coastal states' total, and inside it the band. Equal Earth projection: areas on the page are proportional to areas on the globe, and Antarctica is off the frame but inside the land figure. Coastline SRTM15+ v2.7 (GMT earth_relief 02m), borders Natural Earth 10m, people and output on SSP2 from Murakami and Yamagata (2019), split inside each country on the 1990–2015 grids of Kummu, Taka and Guillaume (2018).

Figure 4. The Indo-Pacific coast. The black diagonals are land within 300 kilometers of the Indo-Pacific; the mesh is the coastal states of that rim taken whole, Russia apart. The bars are a projection for 2050: for population and for output, each shows the coastal states' share of the world in that year and, inside it, the band's share. The projection is equal-area, so area on the page is proportional to area on the earth.

and the legal system each raise a barrier of their own. The wealth produced reaches households only thinly, so labor supply exceeds demand, youth unemployment stays high, births fall and demand shrinks. India, the third candidate, so far belongs to neither.

The division of labor wins because it has become harder to break than it was. The division of labor of the 1990s was an arbitrage on wage differences, so when wages moved the arrangement moved with them. What sets it now is sunlight, land, and a manufacturing lineage accumulated over forty years, and none of those moves as easily as capital or labor. To the extent that comparative advantage is fixed in geography and in time, the gap between the side that divides the terms among itself and the side that holds all of them within one country widens with time. So America, the UAE and Japan-Korea-Taiwan should be bound together more closely than they are. The bits are in America, the physical side is in Japan-Korea-Taiwan, and energy, large capital and the speed of decision are in the UAE. The three already complement one another, but what joins them stops at the level of individual transactions. Put the three under one design and the five terms stand together for the first time. Doing so also makes the cluster, rather than the country, the unit of integration.

One thing could loosen the sea's hold on this map: computation in orbit around the earth. It beats the ground for two reasons, that sunlight arrives unobstructed and that there is no queue for grid connection. No law of physics stands in the way, and the binding engineering constraint is getting rid of heat. Extrapolate today's solid radiators and cooling one gigawatt of computation takes about a hundred thousand tonnes; a liquid-droplet radiator is reported to shed the same gigawatt in about

[43] Inbound direct investment on a realized basis fell 27.1 percent year on year in 2024 to 826.2 billion yuan on the Ministry of Commerce's compilation, while direct investment liabilities on a balance-of-payments basis stayed a net inflow in both 2023 and 2024, the 2023 inflow being the smallest since 1993. The cause cannot be attributed to economic security alone; the contributions of the property market and of interest differentials are not separated out.

2,200.[44] The gap between a hundred thousand tonnes and 2,200 is what separates an impossible project from a possible one; cost is secondary. And at present the lowest cost of access to orbit sits inside the American techno-bloc, so the shift would strengthen that bloc rather than threaten it.

10 What to invest in

The investments map onto the terms of the equation. Because the equation is a product, capital does most where the lowest term is. There is a second test as well: the only returns that can be relied on are the ones that do not depend on head-count.

Raising the effective abundance of energy

- Fusion and SMRs. The contest turns on how many copies of one design get built, so what is worth buying is less the demonstration reactor than the production line and the supply chain that repeat one design after it. The effect on the price of energy arrives in the 2040s, and only capital that can wait that long belongs here.
- Grid infrastructure and heavy electrical equipment. However much generation is added, nothing moves unless a new load can be connected. Three things are buyable: manufacturing capacity for transformers and switchgear, construction capacity for transmission lines and substations, and land already connected to the grid.

Raising labor input

- Physical AI. Humanoids and automated plant. Unit cost is already starting to ride Wright's law while the demonstration of reliability lags, and the two advance separately. None of it can be planned on software's clock. Factories, parts, regulation, accidents and repairs all move in the time of the physical world.
- The leading-edge physical supply chain. Materials, semiconductor production equipment, lithography, leading-edge logic, high-speed memory, industrial robots, precision motors. When volume production begins, what runs out first is not the finished machine but the capacity to make this layer. The region where those layers sit closest together is Japan-Korea-Taiwan, and Japan-Korea-Taiwan is structurally short of energy, so investment in this layer is worth less unless it sits inside the same design as fusion and the grid, because power becomes its constraint.
- Implementation design for the Baumol sectors and housing. Housing, education, childcare and healthcare. Physical AI affects population only when it reaches those four. Building the machines, and designing where in a hospital, a school, a nursery or the production of housing to apply them, are two different investments, and the second is where capital is thinnest.

Building institutions

- The institutions of distribution. Capital accounts at birth, citizen-stake sovereign funds, and universal stakes in automation capital. This is an engineering term rather than an ethical one: it decides whether a cluster is still working twenty years from now.
- The speed of connection and permitting. How good institutions are can be read off how long the procedures take. Britain in the 1840s held two things at once: incorporation by registration alone, and a statute that standardized the procedure

[44] Recomputed by the author from the Stefan-Boltzmann law rather than quoted from industry material. The International Space Station rejects 70 kilowatts across something close to 422 square meters; NASA gives the 70 kilowatts as a design capability and never gives the radiating area, and the 422 comes from the industry analysis cited next. The 450 watts per kilogram for a liquid-droplet radiator is from that analysis, SpaceComputer, "Cooling for Orbital Compute: A Landscape Analysis." It sits at the pessimistic end of NASA's own design studies of the concept, which put the heat rejected per kilogram of radiator at 431 to 2,141 watts (Karl K. Knapp, *Lightweight Moving Radiators for Heat Rejection in Space*, NASA-CR-161943, 1981, Tables I and III).

for compulsory purchase. The interconnection queue, environmental assessment and building permits sit in that position now. The procedure can be shortened in months. The transformer's delivery time can only be shortened in years.

◎ The infrastructure of measuring use-value. Three things are buyable: the technology that records the fact that a beneficiary felt a benefit, the contractual machinery that turns that record into payment, and the machinery that establishes the record was not faked. Without it, neither distribution nor co-creation has a unit to settle in.

Outside the equation

◎ Infrastructure outside the techno-blocs. For the half-century from 2035, head-count is rising mainly in the part of the world that now belongs to neither techno-bloc. The twentieth-century sequence in which cheap labor draws in manufacturing will not be replayed, so that head-count does not turn into purchasing power on its own. Investing to build the ability to buy and investing to sell are two different things: selling into a market where income has not yet arrived turns the sale into a loan.

◎ Orbit. Semiconductors that run hot, and a form of radiator to replace the solid panel, are worth buying. This is the rival to fusion and the grid, and at the same time the insurance on them.

One caution runs through all ten. Most of the returns named here exist because the constraint has not yet come off. Grid-connected land is dear because the procedure is slow, and faster permitting and connection thin it. Ground-based energy infrastructure is dear because orbit has not yet opened, and once orbit opens, the scarcity rent that made it dear goes quickly.

11 Eleven predictions, and what would show them wrong

Writing the predictions down is not enough. Whatever would make me admit one of them was wrong has to be written to the same precision. The table exists not to protect the predictions but to make sure a wrong one can be thrown away without argument.

	Prediction	Observe this and the prediction is lost
I	Energy	By the end of 2035, no year has appeared, in any major economic bloc, in which new capital investment in nuclear exceeded new capital investment in renewables
I′	Fusion (power to the grid in the 2030s, cost at the level of existing sources in the 2040s)	No fusion plant has supplied power to a commercial grid on a sustained basis by the end of 2040, or the levelized cost of fusion reaches the level of existing sources by the end of 2040
II	The grid	Regions that rank at the top on effective industrial electricity price, reliability and interconnection lead time are not chosen, by 2035, as sites for gigawatt-scale new computing and manufacturing
III	Physical AI	By the end of 2032, no producer has reached an annual output of 10,000 humanoids
IV	China's births	Annual births in 2032 exceed 7 million
IV′	Reversal in fertility	A significant reversal is observed in a country centered on cash transfers before it appears in one that lowered the effective prices of energy and housing

V	Labor	The fall in youth employment relative to other age groups, in the occupations with the highest measured AI exposure, levels off in the second half of the 2020s and does not widen thereafter
V′	Retraining	A large-scale retraining program restores the employment rate of the target cohort to a statistically significant degree
VI	Politics	By the end of 2040, a wealth tax or a top marginal rate at the 70 percent level has not appeared in the platforms of parties in government in two or more major advanced economies
VI′	Mobility of the tax base	In a country that implements a wealth tax or a top marginal rate at the 70 percent level, neither a net outflow of top taxpayers nor a shrinkage of the reported base of that tax, relative to independent estimates of top wealth, is observed five years after it takes effect
VII	Distribution	By the end of 2040, none of capital accounts opened at birth, funds held as a national stake, or universal stakes in automation capital has been institutionalized as mainstream policy in a major country
VIII	Demand	By the end of 2040, a shortage of labor still ranks above a shortage of demand in the policy documents of major countries
IX	Measurement	By the end of 2040, not one commercial case has appeared in which payment is settled on a direct measurement of use-value
IX′	Units in energy	By the end of 2040, no commercial contract outside a regulated market settles consideration directly in units of energy consumed or in emission allowances
X	Unit of value creation	By the end of 2040, the average headcount of companies in software and AI has not fallen
XI	Clusters	By the end of 2040, neither Texas nor the UAE is among the leading sites for gigawatt-scale computing and manufacturing
XI′	Orbit	No orbital demonstration at the 50-kilowatt scale by the end of 2032
Zero	The framework itself	Within a window of ten to thirty years, the quality of institutions is shown to move systematically with resource endowment. In that case, the first term of the equation is not an independent term

When one of them is falsified, what breaks in the skeleton? If energy is lost, the timeline slides backward but the identification of the constraint survives. If the grid is lost, the identification itself, that the constraint sits at connection, is wrong. If China's births are lost, the largest case for the diagnosis is wrong. If labor is lost, the break between task and income, the thing that calls for the distribution of ownership, is not yet happening. If distribution is lost, the design this thesis calls necessary will have proved politically impossible. The one whose failure would reach furthest is China's births. The one that most needs to hold is distribution. Some of these will be wrong, and which ones fail, and how, will carry more information than the ones that hold.

12 The institution to build

What I am proposing is a resource integrator, a body that pulls the capital, technology, sites and permits held separately by actors in the United States, the UAE and Japan-Korea-Taiwan into one undertaking: a platform for co-creation that sits between a fund, a studio and a consortium.

As a venture fund, it raises capital and puts it to work. As a studio, it names, program by program, the theme to be taken on, and integrates the resources of the actors who gather under that theme to co-create concepts that do not yet exist. As a consortium, it brings the energy, compute and regulatory parties into one contract. A body that does all three at once is what this thesis means by an ecosystem.

The studio works like this. It sets out one vision of a size no one reaches alone, such as building one fully automated manufacturing plant all the way through. The actors

holding the component technologies that vision requires come together of their own accord to create value, combine their ideas, and conceive, build, run and test a range of products. The membership changes from one program to the next. Putting people from different disciplines in one room and having them design together is not itself new. What changes here is the exit: it does not stop at a proposed concept but attaches gigawatts, compute and a site and carries the work through to the real thing.

The capital starts with the sovereign funds of the UAE, which are working to build domestic industry. It goes into the bits side, the foundation models and AI basic research of the San Francisco Bay Area, Texas and BosWash, and into the physical side, the actors around the leading-edge physical supply chain and the next generation of nuclear reactors in Japan-Korea-Taiwan. The companies invested in then choose, as the ground for their physical experiments, the UAE and Texas, where energy is cheap, a site can be had, permitting is fast and a regulatory sandbox is in place. That is how the UAE and Texas come to be among the leading places for physical implementation.

About the author

Hi, I'm Arata Hoshino, a strategic and venture designer based in the United Arab Emirates. Although raised in Japan, I spent my teens reading about Silicon Valley technology, longtermism, techno-libertarianism, and the ideas of Buckminster Fuller.

I did BEng Mechanical and Energy Engineering at Waseda University; MSc Media Design at Keio University, with research work in service design and venture design; CEMS MIM (MSc International Management) at HEC Paris and the National University of Singapore. After that, I worked in VC and strategy consulting in San Francisco and Singapore. From 2024 to 2026, I advised UAE ministries on economic strategy at a public policy and AI economics advisory firm originating from the Harvard and OECD communities.

This briefing is the short form of *Far Side of the Constraint*, which carries the whole argument and the evidence behind it: nine chapters, 28 figures and more than 120 pages on one basic question, what decides how much a civilization can actually build. It argues that the answer changed over the last fifty years almost without anyone noticing, and follows the consequences through energy, population, value, industrial clusters and investment. It is written for readers who want to check the reasoning for themselves: researchers, the designers of policy, deep tech founders, and investors who build, govern and allocate capital on long horizons. It is a few evenings' reading rather than one, and it ends by stating what would falsify the hypothesis and what it does not claim to prove. If you would like to read it, or to argue with it, contact me.

The thesis, the essays that grew out of it, and the notes that explain it are collected at aratahoshino.com. You can reach me by email, on LinkedIn or on X. The scope of each project is in my portfolio (PDF), and where I stand is in my vision paper (PDF).

Let us meet and talk.
