

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 those who design energy, industrial and population policy. It is a hypothesis that analyzes, conceives and designs on the far side of the present constraint rather than inside it, and it is meant to be used, still as a hypothesis, as a compass for investment and strategy.

Abstract

The argument of this thesis is that contemporary civilization faces one structural constraint, a stalled capacity to build, defined by three deficits: (i) the absence of effective energy abundance, (ii) a shortfall in labor input from the lag in automation, and (iii) the absence of institutions that call in capital and talent. Economic stagnation and population collapse are not the cause, only the symptoms the constraint produced.

The argument runs from the past industrial revolutions, through the diagnosis of the present, to the design of distribution and of value. The industrial revolutions are re-read on two axes, the efficiency of placing matter and the efficiency of prediction, which is what makes the present asymmetry visible. The supply constraint that follows suppresses, at the macro layer, the capacity to conceive large new investment; at the level of the household it raises the prices of exactly those goods automation never reached, and so suppresses population as a second Malthusian trap. Working from that diagnosis, this thesis proposes, in place of allocating income on contribution to tasks, a distribution of ownership built on capital accounts opened at birth. And in place of GDP, on the foundation of service-dominant logic, it moves the center of value to use-value, and sets out the conditions under which it could be measured.

On the far side of the constraint, when energy can be used in abundance and automation arrives at an accelerating pace under the right institutions, the capacity to conceive returns. People are allotted capital before they ever work, and can devote themselves to the creation of value in its proper sense rather than to productivity. The macro problems we now fear, population collapse among them, dissolve as a by-product of the supply constraint being lifted and of a new mechanism of distribution, without waiting on symptomatic treatment, and demand and supply enter a spiral in which each lifts the other.

Finally, the world's nine clusters are analyzed on the five terms the three variables decompose into, and the futures of the American techno-bloc, which runs on a division of labor, and the Chinese techno-bloc, which runs on vertical integration, are forecast. The case is made for deepening the division of labor between the United States, the UAE and Northeast Asia, and for an ecosystem that integrates the actors of the three across their regions. The last chapter sets out what to invest in, term by term, and each of the eleven predictions carries what would have to happen for it to be shown wrong.

A note on method

This thesis project started from numerous notes, written in Japanese between 2018 and 2026. The content and arguments are based on my own ideas. Anthropic's Claude was used to turn those Japanese notes into an integrated English structure, and to create the figures from publicly available data I had assembled. I am responsible for any errors in this thesis.

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, not on a state plan but on private judgment alone.¹ The same country has now spent seventeen years and 46.8 billion pounds without opening a single mile of high-speed line.² Vision is not what changed. The structure of abundance disappeared.

Three terms multiplied together settle what a civilization can newly build in the physical world: institutions, the effective abundance of energy, and labor input. If any one of the three is near zero, the output is near zero however high the others are. Labor input splits further into deep tech and large capital. Effective abundance is not endowment, but whether a society has access to energy that is effectively free, and whether that cheapness is believed to hold for a good part of the life of whatever is about to be built.

The equation does not settle the wealth a country already holds, but the amount it can still build. And the constraint binds not only what gets built. Options are not considered and rejected, but simply stop being proposed. Where a premise is known to be fixed rather than variable, no business plan is written on that premise, nothing reaches an investment committee, and nothing is conceived in the first place. The range of ventures we currently feel to be realistic, and our common sense about the scale of investment, were both built on the assumption that the three terms sit where they now sit. The constraint binds what can be conceived.

Two societies have held all three at once: Britain in the 1840s and the UAE today. In the UAE energy is cheap three times over, through historically subsidized domestic gas, the 5.6 gigawatts of Barakah, and Al Dhafra solar, which set what was then 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, was effectively unlimited. The twenty-first-century UAE is pouring into computing substrate what nineteenth-century Britain poured into railways. Windows also close. Japan's postwar high growth stood on cheap Middle Eastern crude; the window closed in 1973 on a slowdown already under way in the domestic structure, and real growth halved from an annual 9.1 percent to 4.2.⁴

The abundance of 1840s Britain rested in part on child labor and colonies, human costs that must never be repeated, and the labor cost of the present-day UAE stands on a global wage gap. But the combination of energy abundance, fully automated manufacture and transport, and correct institutions makes those three conditions obtainable, for the first time in history, without human sacrifice, anywhere. The abundance the nineteenth century bought with human pain, the twenty-first buys back with machines. That is the civilizational meaning of committing capital to physical AI and energy.

02 Two convergence costs

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; 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; communication, computation, science, education and AI are all special cases of it. Both are ways of lowering the cost of turning an intention into an outcome, one by changing the physical world and the other by knowing what to do.

¹ Odlyzko gives investment in railway construction rising from 13 million pounds in 1845 to 30 million in 1846 and a peak of 44 million in 1847, "almost 8% of British GDP, and twice the military budget of that year" (Andrew Odlyzko, "The collapse of the Railway Mania," rev. 2011). The finer series that circulates, and the 7.3 percent computed against a factor-cost GDP of 604 million, are not in his paper. These are 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, with Ireland a further 537 (B. R. Mitchell, *British Historical Statistics*, 1988, Table 5, p.541, which could not be obtained; in the 1962 predecessor the Great Britain series begins in 1848 and the 1840 figure is a note quoting Lewin, described there as probably United Kingdom). The series counts route length once whatever the number of tracks.

² Cumulative spending was 46.8 billion pounds at the end of March 2026, nominal, of which 44.2 billion is the current programme and 2.6 billion the cancelled 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 design capacity, and with all four in commercial operation from September 2024 it generated about 40 terawatt-hours in the following twelve months, which is up to 25 percent of national electricity demand on an annual basis rather than a share of instantaneous output (Emirates Nuclear Energy Corporation). Al Dhafra solar is 2 gigawatts of nameplate capacity, about 1.5 gigawatts on an alternating-current basis, and its April 2020 award price of 1.35 cents per kilowatt-hour was the lowest announced anywhere at the time, revised to 1.32 cents at financial close (Abu Dhabi National Energy Company, press release of 22 December 2020). The IMF's Article IV report put implicit energy subsidies at 5.7 percent of GDP in 2011, falling to an expected 4.6 percent in 2015, of which natural gas was 3.0

The two compete over the same problem, and the side with the lower convergence cost wins. How the problem splits between them can be measured before seeing who wins. Decompose the problem into steps and count the mass that must be moved at each step. For shopping, search, comparison, ordering and payment all move zero mass. The only step that does not is the last leg, from warehouse to doorstep. So this problem could be read in advance as one that would mostly move to direction 2. That Amazon beat the shopping robot people used to imagine was not a matter of the superiority of one vision.

Capital and talent flowed one-sidedly into direction 2 for thirty years because that side needed less energy. Direction 2 replicates at zero marginal cost, runs with low friction, and needs little energy. Only the last of the three varies with the era, and in a constrained world capital flows toward the looser constraint. Then, in the 2020s, direction 2 collided head-on with the wall of energy, and improved efficiency could not avert it. A fall in convergence cost unlocks applications that were not previously viable, so 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 has been falling rapidly. Cheaper computation was not saved, but used more widely, more often, and for heavier purposes.

The movement of capital speaks more plainly than forecasts. In August 2025, the Palisades plant in Michigan was returned by the regulator from the decommissioning process to an operating licence, the first time such a restart has been approved, which is a regulatory category and not the resumption of generation.⁶ The structural change this thesis depends on is that computing demand has made investment in energy supply politically unavoidable. The commitment of capital to physical AI and energy is not the next trend but capital's historical return to the side that was left behind.

03 A scarcity chosen

The slowing of growth across the advanced economies is not the middle-income trap that stalls developing countries but an advanced-country trap. Of the three terms, the one that broke was energy.

There are three exhibits. First, on the annual series the EIA publishes, which begins in 1949, primary energy consumption per capita in the United States climbed from 207 million Btu a head to a peak of 350 in 1978 and 1979, and has since come down to 282.⁷ A climb that had run through the whole postwar period stopped in that one decade and has not resumed, and the level is now close to a fifth below the peak. Since efficiency gains have repeatedly raised total consumption rather than lowered it, this reads not as the fruit of efficiency but as the halt of expansion. Second, in Robert Gordon's aggregation, total factor productivity in the United States grew at roughly 1.9 percent a year from 1920 to 1970 and after 1970 fell to about a third of that rate, with no recovery since. In the model with which Ayres and Warr asked how much of growth the useful work extracted from energy can explain, the fit to American output holds from 1900 until 1973 and 1974 and then begins to fall short.⁸ Third, all 41 reactors ordered in the United States after 1973 were canceled.⁹

This was not a consequence of physical law. 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, while 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 repeatedly. The same technology escalated in one country and fell in

points (IMF Country Report No. 15/220).

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).

⁵ IEA, *Key Questions on Energy and AI*, 16 April 2026 (data center electricity demand up 17 percent in 2025, AI-specific up 50 percent; 485 TWh in 2025 rising to about 950 TWh in 2030 on the central projection, both for data centers as a whole rather than for AI alone). The IEA states that energy per AI task is falling by at least an order of magnitude a year; it publishes single estimates measured under test conditions but no consistent annual series for the level.

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). "Operating" here is a regulatory category, not a resumption of generation. 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. This thesis uses only 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).

another, and the difference was regulation and politics.¹⁰ The 1970s did not simply happen, but were chosen, and the problem is that 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 here is that rebuilding; it was a scarcity chosen, not a scarcity given.

The main exit the advanced countries chose in that half-century was a source that swings against human need. 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 it generates, the more everybody else is generating in the same hours, so it destroys the value of the very electricity it produces. The number of hours in which German wholesale power prices were negative reached 573 in 2025, a record for the second year running.¹¹ A society that handed the baseload of its civilization to a source that belongs on top of firm power as a supplement distributed the cost of doing so, in the form of the electricity bill, to households. 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 were printed on the ballot paper, it would be a democratic choice. All that was missing was the awareness of having chosen.

04 The household Malthus

The price of a good or a service is set, roughly, by energy cost, manufacturing cost, transport cost, computing cost, and the labor cost that 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 items that make up raising one child are almost all 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 is the act of buying, twenty years at a time, precisely the goods that automation has reached least.

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.


¹⁰ Jessica R. Lovering, Arthur Yip & Ted Nordhaus, "Historical construction costs of global nuclear power reactors," *Energy Policy* 91, 2016, 371–382. France connected 58 units to the grid between Fessenheim 1 in 1977 and Civaux 2 in 1999, 52 of them by 1991.

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, which is 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.

¹³ Bureau of Labor Statistics CPI series, retrieved by the author through the API and computed directly. From January 2000 to January 2024: hospital services +248.9 percent, college tuition +184.6

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. Children are bought from the upper bundle.

So falling fertility is not a cultural phenomenon but a price phenomenon. 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 paying out, on a budget line that grew more than fortyfold, and the price structure itself did not move. Fertility rates are now separated not by income but by the cost of combining career and family.¹⁵ Values are not a primary cause, but appear two or three steps down from a structure that sits below them. Energy and technological structure settle the price system, and the price system settles the household's budget constraint and its anxiety about the future. Under that constraint, households weight career over children as a rational principle of action, and the principle later justifies itself as a value. To detect a change in a downstream variable and call it a primary cause is close to calling the movement of a thermometer needle the cause of the heat.

The site 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 it will not reverse is already settled 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 Population is the variable that is fixed before any other, and the one that moves last.

05 The cascade

On a week away, with only 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.

percent, day care and nursery school +134.3 percent. The figures in the text are more conservative than the measurements. On the falling side, televisions are down about 97 percent over the same period and toys and computer software about 70 percent each. These are quality-adjusted indices and the electronics series carry hedonic adjustment, so a 97 percent fall is the price of a constant quality of television rather than 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: on national funds, 8.93 trillion won across the first basic plan, 27.96 across the second and 118.75 across the third, for 155.63 trillion over 2006 to 2020, with 42.90 budgeted for 2021. These are budgeted amounts rather than settled expenditure, and the low-fertility side only. The 378.1 trillion won carried earlier is not in that report. What counts as a low-fertility measure widens over the period until youth employment and housing are 61 percent of the 2021 total, so it is not a single series on one definition. TFR is Statistics Korea, 2023 birth statistics (0.72 nationally, 0.55 in Seoul). It has since recovered to 0.75 in 2024 and a provisional 0.80 in 2025; this thesis does not use the reversal as evidence, because a rebound cannot yet be separated from it.

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.

Latent business models branch like a tree, and to stop short because a constraint is in the way may be to cut away, wholesale, every branch that would have flowered further on. On energy, our civilization is that traveler with 2GB.

So what happens when the cost of energy collapses is not a discount on existing demand but an explosion of demand. One of the goods whose cost humanity has driven down for longest is artificial light. On Nordhaus's measurement, the real price of a thousand lumen-hours fell by 99.97 percent between 1800 and 1992.¹⁷ So when light became three thousand times cheaper, did humanity spend less on light? 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, humanity consumed more light. It invented the city at night, the twenty-four-hour factory, the screen that never sleeps. A share of income that stays in the same band while the price collapses is what a rebound close to complete would look like, and that is the reading this thesis takes.¹⁸

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 per cent 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 per cent of GDP across seventeen observations with a best fit of 0.72. The three points on the right are the empirical settings the paper reports, spanning tallow candles to grid electricity.

Light is the clearest historical case, but the mechanism is not peculiar to lighting. What gets unlocked can be drawn concretely from the cost structure. Between a quarter to 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 that Jones attributes to Musk.²⁰ Complete the third order of magnitude and orbit becomes an industrial location.

The order of supply runs from fusion at the top, through conventional large reactors and small modular reactors built in series in a factory, with fossil fuels as a bridge and renewables riding on top of firm power as a supplement. The price is set not by discovery but by volume. The cost curve follows Wright's law, by which unit cost falls at a constant rate for every doubling 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 at converting 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. Yet between the ignition of the first machine and a tenfold reduction in cost per unit lie several decades and several hundred reactors. Prices genuinely move in the 2040s.

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 outturn.

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. The underlying data are Tsao & Waide, *LEUKOS* 6(4), 2010, 259–281, whose abstract gives the span as three centuries, six continents, five lighting technologies and five orders of magnitude, and which gives the share as 0.39 to 1.30 percent of GDP across seventeen observations with a best fit of 0.72. Neither paper puts a number on the rebound, and the step from a stable expenditure share to a rebound close to complete is this thesis's own. It is, however, measured elsewhere: Fouquet & Pearson (*EEEP* 1(1), 2012) estimate the rebound directly for British lighting and find that nineteenth-century efficiency gains produced backfire, energy used for light rising rather than falling.

¹⁹ On 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. For indoor vertical farms, a review across many facilities (Pennisi and colleagues, *Agronomy for Sustainable Development* 45, 2025) gives lighting at 42 to 80 percent of energy use and climate control at 16 to 43 percent, and the ranking of the two swaps with facility type and climate. The share of operating cost is where the published figures pull apart, and the difference is one of premises: a paper design of a tower on German commercial tariffs, with capital excluded from the denominator, puts energy near seven tenths (Banerjee & Adenauer, 2014), while a survey of more than 150

06 What sets the speed is not generation

Generation itself has grown exponentially. Cumulative installed solar capacity climbed six orders of magnitude, from 0.54 megawatts in 1975 to 1.87 million megawatts in 2024, and China accumulated 1,840 gigawatts of wind and solar alone by the end of 2025, 47 percent of its total installed capacity.²² That does not make the growth in delivered energy exponential.

The constraint sits between the generator and the load: the queue for new grid connections, manufacturing capacity for transformers and switchgear, land acquisition and permitting for transmission lines, substation construction. The median period from application to entry into operation, for projects entering operation in the United States in 2025, was 61 months, and of the projects applied for between 2000 and 2020 only 13 percent by capacity reached operation at all.²³ A new interconnection right is no longer bought with time, but with probability. The delivery time for large transformers is now two to three years, with generator step-up transformers at 143 weeks and power transformers at 128.²⁴

Every one of these compresses only in years, however many factories you add, and all of them tightened across the major markets at the same time from the early 2020s. Two implications follow. Prices in the transition rise rather than fall, because demand explodes while supply is held. And the scarce asset is not the generating plant but the site already connected to the grid, and the position in the queue. Whether growth is exponential at all is settled at that connection. And since neither electricity nor land can be exported, this is the one term of the equation that cannot be divided among partners, which makes the choice of site itself the investment decision.

07 The valley of demand

Supply grows at once; demand grows only at the speed of population. For an economy to work, it needs at least four head-counts: labor supply, demand, risk-takers and taxpayers. Machines can replace only the first. If nobody wants the thing, then however much you make, it is only inventory. How many people will put their own lives into a business nobody has tried is settled not by productivity but by the size of the population. You can tax the wealth machines generate, though a human community is what carries that taxation, decides its distribution and maintains its institutions.

The four do not come back at the same speed. The time it takes one human being to join each can be read off the National Transfer Accounts.²⁵ Demand is same-day: from the moment of birth a person is a consumer. Labor supply and taxpayers take decades. On that profile a person's labor income does not exceed their own consumption until 30, and for all of those years they are on the receiving side rather than the paying side. Risk-takers take longer still: the mean age of founders of firms in the top 0.1 percent by growth is 45.0, and the probability of success goes on rising to nearly 60.²⁶

Even on an extremely optimistic assumption, the valley lasts half a century. Suppose the constraint comes off completely in 2035, and that at that instant the total fertility rate of high-income countries jumps in a single step from 1.42 to 2.1 and stays there. This cannot happen in reality, yet since no faster reversal is possible either, the valley this calculation produces is a lower bound.²⁷

operating North American farms gives 12 percent for small sites and 25 for large (Agrilyst, 2017). Most of the gap is the cost of labor 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* (of applications filed between 2000 and 2020, 13 percent by capacity reached operation and 75 percent were withdrawn). These are generation-side statistics; the load-side queue is not covered.

²⁴ Wood Mackenzie, *Making the connection: meeting the electric T&D supply chain challenge, Executive Summary*, September 2025, with American Clean Power, slide 5: for the second quarter of 2025, generator step-up transformers at 143 weeks after receipt of order, power transformers at 128 and switchgear at 44. The Department of Energy's *Large Power Transformer Resilience: Report to Congress* (July 2024) gives the same picture in months, at 36 typically and up to 60. The figures are for the American market.

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 falls furthest.

The investment implication follows. Through the half-century from 2035 to roughly 2085, what can be relied on is returns that do not scale with head-count: stakes in automation capital itself, and the infrastructure of energy and computation. Businesses whose returns are proportional to the volume of demand sit in a structural headwind for that half-century. In a world short of demand, the strong position belongs not to the side that can make but to the side that can buy. And through that half-century population is still rising mainly in East, West, Central and North Africa, the Levant, Pakistan, and parts of Central Asia and Central America, where purchasing power per head is still low. But demand is population multiplied by purchasing power per head, and what that part of the world holds now is future demand, not present demand. So for the side holding plant it cannot sell, building up the other side's ability to buy is not charity but the work of turning its own plant from a liability back into an asset.

08 The distribution of ownership

This thesis is optimistic about supply and pessimistic about distribution. To put it precisely, the prediction here is that redistribution by taxation will fail once, somewhere in the 2030s.

Modern economies made the task the instrument by which wealth is distributed. Work was broken into discrete tasks, and income has been allocated through the performance of those tasks. Do X and you are paid Y. Reduced to its essentials, distribution in the modern economy runs on that one line, and that line began to fail before AI. The labor share has fallen significantly since the early 1980s across the majority of countries and industries, and about half of that fall is explained by machines and computers getting relatively cheaper and firms replacing people with them. The displacement effect, by which automation hands a

²⁵ National Transfer Accounts Project, per-capita profiles by age; this thesis uses the 2011 United States profile. 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. The ages at which net payments to the public sector turn positive, which this thesis carried earlier, could not be traced to a published source and have been withdrawn. 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 base population is World Bank data for high-income countries by five-year age band and sex, 2025, interpolated to single years and run through a cohort component projection. Age-specific fertility is a skew-normal distribution with a mean of 29.0, mortality is a Gompertz-Makeham single-year survival schedule calibrated to the region's life expectancy and then held fixed, and migration is set to zero.

human task over to a machine, always lowers the labor share; what offsets it is the reinstatement effect, in which new tasks are created for humans. In the United States of recent years, displacement accelerated and reinstatement weakened.²⁸ And the break is an observation of the 2020s rather than a prophecy about the 2030s. In the occupations AI reaches most, employment of young workers fell relative to other age groups while employment of experienced workers held steady; the fall came through a halt in hiring rather than through layoffs, and the gap stands at 19 percent as of June 2026.²⁹ It is breaking at the entrance, at the young end.

Politics always reacts. In November 2025, candidates standing as democratic socialists won the mayoralties of New York and Seattle, and Gallup's survey that year put favorable views of capitalism down to 54 percent.³⁰ But the remedies this wave puts forward will fail, and the mechanism of failure is not the old point about high tax rates blunting the will to work but capital flight. Around 1990, twelve OECD countries levied a tax on net wealth; three remain.³¹ Where capital and people can move, taxing wealth country by country has not historically been sustained. On cash transfers, we now have the best experimental data there is: a thousand people given a thousand dollars a month for three years, against two thousand on fifty, worked 1.3 hours a week less, their employment rate fell 4.2 percentage points, and health indicators showed no measurable improvement.³² Cash buys autonomy and immediate peace of mind. A cheque does not change who owns the capital.

This thesis puts forward instead a design that moves the ground of distribution from tasks to the ownership of capital. As the labor share falls structurally, the only way for households to survive is to move over to the capital side. The institutional seedlings are already in the ground. Under a law enacted in the United States in 2025, the federal government pays 1,000 dollars into an account for every child born between 2025 and 2028, invested in a low-cost fund tracking a broad index of American equities. On the Alaska Permanent Fund Dividend, the negative effect on aggregate employment was found to be zero. Norway's Government Pension Fund Global is collective capital of roughly 380,000 dollars per citizen.³³ Broad middle classes have repeatedly been made by turning citizens into owners: the United States moved about 10 percent of its national territory into private hands through the Homestead Act, and postwar Japan moved 1.74 million hectares to 4.75 million tenant households and cut the tenanted share of farmland from 45.9 percent to 9.9.³⁴ The nineteenth century distributed land. The twentieth distributed home ownership and pensions. The twenty-first has to distribute a stake in automated capital.

The same technology works from below as well. 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 itself. However, if the threshold falls without any design for dispersing ownership, then business owners without capital multiply rather than capitalists. And there is a temporal asymmetry in the order. Institutions of distribution can be made by legislation, so once the decision is taken they move within a few years. Population does not. A child born now is a consumer from the first day, yet does not reach risk-taking or the tax rolls for another two decades. To institutionalize distribution in the 2030s is to buy the demand, the challengers and the tax revenue of the 2050s now. Institutions can solve the problem of income, not the problem of time. Which is why there is a need to hurry.

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 (50 to 70 percent of the change in the American wage structure). The "about half" decomposition is contested on its sign, not its size: Elsby, Hobijn & Sahin (*BPEA*, Fall 2013) find the coefficient on the relative price of investment goods running the wrong way, and Bergholt, Furlanetto & Maffei-Faccioli (*AJ: Macroeconomics* 14(3), 2022) find automation rather than cheaper machines doing the work.

²⁹ Erik Brynjolfsson, Bharat Chandar & Ruyu Chen, "Canaries in the Coal Mine?" Stanford Digital Economy Lab, 2025, revised 2026. The authors themselves note this is correlation rather than causation, and the figure has moved with the data: 13 percent in the August 2025 version, 16 percent in November 2025, 19 percent in the June 2026 follow-up. The same age gap appears in Danish registers, where 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 and call their result an external validation of the Stanford finding.

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, "Why were most wealth taxes abandoned and is this time different?" *Fiscal Studies* 42(3–4), 2021, 539–563. The fact is not a proof that taxing wealth is wrong, only that it has been hard to sustain under mobility; Zucman's proposal for international coordination is an attempt to change that condition.

³² Eva Vivalt, Elizabeth Rhodes, Alexander W. Bartik, David E. Brookman, Patrick Krause & Sarah Miller, "The Employment Effects of a Guaranteed Income," NBER Working Paper 32719, 2024, revised 2026. The

09 Moving the center of value

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 loving her child are the same act, and the only difference is that one of them carried a price. Adam Smith drew the line. In Book I, Chapter IV of *The Wealth of Nations* he notes the two senses of value, use-value and exchange-value, sets down the paradox that nothing is more useful than water and almost nothing can be bought with it, and then immediately confines everything that follows to exchange-value. The reason was practical, not metaphysical: exchange-value gave a standardized measure. And he called use-value the real value. The paradox was not solved but shelved, and then the shelving was forgotten.³⁶

The size of the blind spot can be measured. In the satellite accounts of the US Bureau of Economic Analysis, unpaid household production is worth twenty-two to twenty-five percent of GDP, about a quarter of which is the service flow from durable goods rather than unpaid time. Averaging the OECD time use surveys over ages 15 to 64, paid work and study come to only 4.6 hours a day, against 3.4 to 3.5 hours of unpaid work and about 4.8 of leisure.³⁷ And in large-scale choice experiments, the median user's compensation for giving up one free platform for a month was about forty-eight dollars, for a good with near-zero contribution to GDP.³⁸ All three ways of measuring point the same direction. The value that market transactions count is a minority of the value we create in a day. This is not a bug in GDP, but the specification, Smith's proxy institutionalized over two hundred and fifty years.

Price does two jobs: it rations scarcity and it measures value. If the monetary institution is to retire, both have to go. The first condition has already been observed in information goods, where marginal cost near zero drives price toward zero and price stops working as a signal of scarcity; the collapse in the cost of energy and automation extends the same mechanics into the material world. The second is not settled by energy. Price as an institution stands on one fact and no other, which is that at the time there was no way to observe use-value, so something observable was put in its place. That premise is the one now breaking. The point of observation moves from the transaction to the body, and from the body to the nervous system. If the very fact that a beneficiary felt a benefit becomes recordable, the reason for keeping a proxy disappears.

The two do not arrive at the same speed. The rationing job goes for reasons of energy, on a scale of decades. The proxy job goes for reasons of measurement and of trust, and that is probably the harder of the two. So the place where money first gives way is not the side that measures value in data but the side of energy. The token that counts AI inference, and the carbon emission allowance: both were designed as units of account and both are converging on currency. Anything issued, exchanged and stored as a unit of a scarce resource is performing the unit-of-account and store-of-value functions of money at that point. For an investor the implication is one line. Once measurement is cheap, the claim can attach to the use-value of the thing made rather than to the organization that made it, and investing becomes the purchase of a claim on use-value that has not yet arrived, rather than the purchase of a share in a company.

10 The center of gravity has gone to sea

The body of the world economy is not the nation but the cluster. Cut the world's land by the continuity of its lights at night and the lumps of the economy come out on their own:

control group received 50 dollars a month, so the estimates are the effect of a further 950; the 4.2 points is on employment, with a standard error of 1.7 against a control-group rate of 63 percent. Partners cut their hours by about as much again. On health, Miller and the same colleagues, NBER Working Paper 32711, 2025. The Finnish experiment likewise raised subjective wellbeing while its employment effect was close to zero, at 1.5 days in the first year with a confidence interval from -2.3 to 5.4. 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: the Treasury pays 1,000 dollars for each United States citizen child born between 1 January 2025 and 31 December 2028, and the investment is restricted to funds tracking a broad index of American equities. On Alaska, Damon Jones & Ioana Marinescu, *AEJ: Economic Policy* 14(2), 2022, 315-340. Norway is 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 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 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 follows Lusch & Vargo, 2014, ch.2.

³⁷ 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; the paper's prose calls it a share of the extended satellite output, but its own

the 40 megaregions Florida and his colleagues extracted this way hold 17.7 percent of the world's population, and out of them come 66 percent of world output and 85.6 percent 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 as of the twenty-first century, it is not that clusters are arranged along the rim of the Indo-Pacific but that the cities along that rim are the ones that became clusters. More than eighty percent of world merchandise trade by volume, and more than seventy percent by value, moves by sea.⁴⁰ The physical substrate treated here, the transformers and turbines and reactors and building materials, is all on the heavy side. Distance from a port is, directly, distance to construction.

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 Kumm, 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 show the coastal states' total in 2050 and, inside it, the band. The projection is equal-area, so area on the page is proportional to area on the earth.

Land within 300 kilometers of the Indo-Pacific, from the Bering Strait round to East Africa, is 12 percent of the world's land. By 2050, 34 percent of the world's people live on it and 39 percent of world output comes off it.⁴¹ Take the littoral states of the Indo-Pacific together and they hold 69 percent of the world's people and produce 66 percent of world GDP at purchasing power parity; in 1980 the figure was 53 percent.⁴² That the merchant cities strung along the rim of the largest ocean on earth grew giant is not a difference of effort, but a condition they were dealt first.

A cluster's standing is settled not by how many terms it has assembled for itself but by whether it can reach the ones it lacks. Seen in this unit the equation splits into five terms: institutions, effective energy abundance, deep tech in bits, deep tech in the physical, and large capital. This thesis calls the bundle of those reaching relations a techno-bloc, and there

table divides by published GDP). Time use is the OECD Time Use Database and Veerle Miranda, OECD Working Papers No. 116, 2011, on the average for 29 countries. The average covers all days including weekends and both employed and non-employed, which is why paid work looks short. The pairing of 4.5 hours of paid work with 4.1 of unpaid work for ages 25 to 64, which this thesis carried earlier, is not in that paper.

³⁸ Erik Brynjolfsson, Avinash Collis & Felix Eggers, "Using massive online choice experiments to measure changes in well-being," *PNAS* 116(15), 2019, 7250-7255.

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. Output is estimated from the intensity of night-time lights and renormalized to national GDP, which is 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. The 2024 and 2025 editions do not state this ratio, so the 2017 edition is cited.

are at present only two that actually have the terms assembled. The American techno-bloc assembles them by division of labor; the Chinese techno-bloc tries to assemble them by vertical integration. The third candidate is India, which so far belongs to neither way of assembling them.

The essence of the American techno-bloc is not that one country holds every term but that a division of labor holds. The Bay Area runs out along deep tech in bits, Japan-Korea-Taiwan along deep tech in the physical, and Texas, BosWash and the UAE along institutions, energy and capital. Not one of them is complete, and put together the gaps disappear. The Chinese advantage is speed of decision and resilience against a cut in supply; the defect is that if any one term is structurally low it becomes the bottleneck for the whole. Institutions are what is actually blocked. Direct investment coming in from abroad has turned down amid economic-security concerns, and residency is not detached from the employer.⁴³ The wealth produced reaches households only thinly, so labor supply exceeds demand and youth unemployment stays high. Births fall with it, and demand shrinks.

That reading holds because the division of labor 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. Sunlight, land, and a manufacturing lineage forty years in the accumulating settle it now, 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. Putting them under one design also means bringing the unit of integration down from the country.

One thing could upset that arrangement from outside: orbit. Computation in orbit 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 2,200.⁴⁴ The distance between a hundred thousand tonnes and 2,200 is not a difference of cost but of possible and impossible. And at present the lowest cost of access to orbit sits inside the American techno-bloc.

11 What to invest in

The places map onto the terms of the equation. Because the equation is a product, capital does most where the lowest term is. A second test: 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 itself 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.

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 Kumm, 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. 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.

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.

⁴⁴ The figures in this passage were 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, under a fifth of the ideal; NASA gives the 70 kilowatts as a design capability and never gives the radiating area, and the 422 is from the same industry analysis. The 450 watts per kilogram for a liquid-droplet radiator is from SpaceComputer, "Cooling for Orbital Compute: A Landscape Analysis," an industry analysis whose primary source is unverified.

- ⊙ 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 much less unless it sits inside the same design as fusion and the grid.
- ⊙ Implementation design for the Baumol sectors and housing. Housing, education, childcare and healthcare. Physical AI tells on 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 the thinner of the two.

Building institutions

- ⊙ The institutions of distribution. Capital accounts at birth, citizen-stake sovereign funds, and universal stakes in automation capital. This is not an ethical term but an engineering 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 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 only 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 stack. For the half-century from 2035, head-count is rising mainly in the part of the world now placed outside the stack. 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 does open, the scarcity rent that made it dear goes quickly.

12 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. This 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	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'	Capital flight	In a country that implements it, neither a net outflow of the wealthy nor a fall in revenue from that tax 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 backwards 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, which is what calls for the distribution of ownership, is not yet happening. If distribution is lost, the design was necessary and yet 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 the way they go wrong is the information.

13 The institution to build

What I am proposing is a resource integrator that connects the terms held separately by actors in the United States, the UAE and Japan-Korea-Taiwan: 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. When one body does all three at once, what results can be called 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 gather of their own accord to create value, lay their ideas over one another, conceive a range of products, build them, run them and test them. The faces change fluidly with the vision. The practice of putting people from different disciplines in one room and having them design together is not itself new. The exit is what changes here, since it does not stop at a proposed concept but attaches gigawatts, compute and a site and carries the work through to the real thing. It is an ideathon run to completion.

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. Those 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. This is how the UAE and Texas come to be among the leading places for physical implementation.

14 Where this reading could be wrong

The heaviest case is that capital markets break first. In Bank of America's fund manager survey of February 2026, the share of respondents saying that companies are overinvesting in plant reached the highest level in the history of the survey.⁴⁵ Let AI investment break badly once and the expansion curves of energy and computing power themselves bend. At that point, the diagnosis of what the constraint is stays correct, and only the date it comes off slips back by a decade or more.

The calculation behind the valley of demand is not a forecast. Migration is set to zero, which works to make the valley in high-income countries look deeper than it is. The timing

⁴⁵ Bank of America Global Research, *Global Fund Manager Survey*, February 2026. 162 respondents with 440 billion dollars under management.

Overinvestment was named by 35 percent, the highest in a series going back to 2005. As the source of a systemic credit event, private equity and private

survives rather than the level, in that the year of the floor does not leave the range 2085 to 2090 even when the life-expectancy assumption is moved four years either way. The same holds for China. The National Bureau of Statistics does not publish a total fertility rate, and the 0.93 used here is a private estimate. The population base is robust, which is to say that every woman who will give birth in 2050 has been born by 2032, and the absolute number of births is not.

Price explains the conditions under which fertility falls; it does not explain the conditions under which it rises. That fertility collapsed from the 2010s in the Nordic countries, which hold the most generous family policy in the world, and that Israel holds close to 2 even among its secular population in a high-price society, both mark the edge of this explanation.⁴⁶ No record exists of a state in which the constraint has come off, so what remains are the episodes in which it loosened, and those are too few to found a law on. Which is why this is written as a conditional prediction rather than as a law. On clusters, geopolitics is not treated here, and shifts in power can overwrite the arrangement. And the argument that use-value becomes measurable is the most distant of all: so long as it leans on a proxy, we may only have replaced one proxy, price, with another, biometric data or some other behavioral proxy, and an institution that pays for measured pleasure creates at the same time the incentive to manufacture it.

credit led at 43 percent, with hyperscaler capex second at 30 percent.

⁴⁶ On the Nordic countries, Hellstrand, Nisén, Miranda, Fallesen, Dommermuth & Myrskylä, *Demography* 58(4), 2021. Israel is the Central Bureau of Statistics, *Statistical Abstract of Israel 2025*, Table 2.28 (TFR of 2.87 across the whole population in 2024, 1.98 among the non-religious). The religiosity categories are self-reported in social surveys.

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.

You can reach me by [email](#). A summary is on [LinkedIn](#), the scope of each project is in my [portfolio \(PDF\)](#), and where I stand is on [vision paper \(PDF\)](#).

This document is the skeleton of a completed thesis of nine chapters, roughly 56,200 words and 28 figures, in English. The full text is not published. Every figure was drawn by the author from public data.