

Everything Is Getting Cheaper

The falling cost of the four inputs of an economy, and what it reorders

Muzamil Hasan

Founder, arc. · Host of *Thought Behind Things* and *Beyond the Bubble*
muzamilhasan.com · hello@muzamilhasan.com

A working paper · 2026

Abstract. Strip an economy to its parts and it runs on four inputs: intelligence, labour, energy, and raw material. For the first time in recorded history, the cost of providing all four is on a downward trajectory at once. The claim is not that things are cheap today; most of these inputs are getting more expensive first, and that expensive phase is part of the mechanism, because rising prices are what pull investment and invention into a scarce input. History shows that when the cost of a fundamental input collapses, growth leaps, work is remade, and economic power moves. The recent decades of slow growth are not evidence against this; they are the below-trend stretch that the collapse in the cost of intelligence is positioned to end. For the emerging world, and for Pakistan specifically, this rebalancing is an opening, but an opening that has to be earned against a hard base rate, not a destiny. For an individual, the move is to participate early: apply the cheapening input to a domain you already understand, and build where building has become cheap, because when building is cheap the advantage shifts to judgment and distribution. Every major claim in this document carries the condition under which it would be wrong.

What is new here, and what is not

Most of the raw material of this argument is not mine, and I will cite whose it is as I go. The abundance tradition, the long-cycle view of economic power, and the economics of induced innovation are old, serious bodies of work. What I am contributing is the coupling. First, a four-factor lens: treating intelligence, labour, energy, and raw material as the complete input stack of an economy and tracking the cost trajectory of each. Second, a measurement rule: track the cost of providing an input, not its market price, because price follows cost down only where competition is free to work. Third, the connection between these mechanics and the macro question of where economic power moves next, argued to the audience with the most at stake and the least written for it. Where I am extending someone else's work, I say so. Where the evidence is against me, I say that too.

How I came to this

A word on where this comes from, because it is not an academic exercise for me and I would rather it not read as one. In 2010, as a teenager in Pakistan with a borrowed camera and no connections, I walked into media at the moment its cost collapsed, and I spent the next fifteen years watching what that collapse did: to me, to the thousands of others who walked through the same door, and to the ones who stood outside it arguing about whether it was real. The podcast that grew out of that work put me across the table from several hundred founders, economists, and scientists, and the pattern in this paper is the thing I could not stop seeing in those conversations. So the framework is assembled from other people's research, which I cite throughout, and the conviction behind it is my own, earned the slow way. I write in the first person

because I mean it in the first person. If any part of this is right, it should be argued with, extended, and used; my details are at the end for exactly that.

1. Everything is getting cheaper

The claim of this document is in its title, so I should state immediately what the claim is and what it is not.

It is not a claim that life is cheap today. Your rent is higher than last year. So are your groceries, your fuel, and your electricity. That experience is real, it is measured, and no honest argument can begin by denying it.

The claim is about trajectory, and about a specific set of costs: the cost of the inputs everything else is produced from. Those costs are heading down, some of them faster than any cost has fallen before. And, which is the part that takes longest to accept, several of them are getting more expensive right now as a direct step on the way down.

Two ideas carry the whole argument, so I will set them out first.

The first idea: the four inputs. Forget countries, currencies, and markets for a moment. Underneath all of it, an economy is four things working together: intelligence, labour, energy, and raw material. The ability to think, the ability to do, the power that runs the doing, and the material you shape. Everything humanity has ever produced is a combination of these four, and whichever one is scarcest sets the ceiling on growth. This is an old intuition in economics. Liebig's law of the minimum said growth is set by the scarcest resource, and Michael Kremer's O-ring theory showed how one weak input can cap the value of all the others.¹ An economist will notice that I have left out two factors the discipline cares about most, capital and institutions, and the omission is deliberate rather than ignorant. I am tracking the physical inputs whose cost is falling. Capital responds to those costs, and institutions decide who captures them. Institutions return in chapter six, where they become the whole point, because cheap inputs do not automatically make a country rich, and the reason is institutional.

The second idea: expensive before cheap. When demand for an input rises faster than supply, its price rises. The rising price does two things at once. It hurts, visibly and immediately. And it redirects money, research, and talent toward producing more of that input, which is what eventually collapses its cost. This is not my mechanism. John Hicks described it in 1932: when a factor of production becomes dear, invention turns to economizing on it.² Daron Acemoglu formalized it as directed technical change, with the caution that the direction depends on how substitutable the inputs are, so the loop is a strong tendency rather than a law.³ And it has been measured. David Popp found that when energy prices rise, energy-saving patents rise behind them, with a long lag and diminishing returns.⁴ Newell, Jaffe, and Stavins found the same, and added the caution that policy and standards did much of the work, so the response is real but not automatic.⁵

Energy is running this loop in public right now. After roughly two decades of flat electricity demand in the developed world, AI computing has made demand surge. United States electricity prices rose 6.9 percent in 2025, about twice inflation, with data centres accounting for roughly forty percent of demand growth.⁶ The response is already visible: a record 3.3 trillion dollars flowed into energy investment in 2025, about 2.2 trillion of it into clean energy and grids, with the International Energy Agency attributing the acceleration substantially to AI demand.⁷ Retired nuclear plants are being restarted under long-term contracts with AI companies, and money has moved into fusion and next-generation storage at a scale the

field has never seen.⁸ Whether all of this converts into cheaper energy is the bet this document makes, and I put a date on it below. But the sequence itself, demand up, price up, investment up, is not a forecast. It is in this year's data.

One more rule, and it matters everywhere in this document. **Track the cost of providing an input, not its market price.** Price is cost plus margin. A high margin does not contradict a falling cost; it advertises it, because fat margins are what pull competitors in, and competition is what pushes price down toward cost. The cost of launching a kilogram to orbit is a different number from what a launch company charges. The cost of generating a unit of electricity is a different number from your tariff. Where entry is free, price follows cost down with a lag. Where entry is blocked, by monopoly, cartel, or the political structures I discuss in chapter six, price can stay high while cost collapses underneath it. When that happens, the gap is not a refutation of this thesis. It is a measurement of how much is being extracted between the falling cost and you.

So the full claim, stated precisely: the cost of providing intelligence, energy, labour, and raw material is on a long downward trajectory; two of the four are falling now, one is turning, one is further out; the expensive phase several of them are passing through is a step in the mechanism, not evidence against it; and the consequences, for work, for countries, and for individuals, follow from the trajectory rather than from today's price.

None of this would be believable if it were happening for the first time. It is not, and that is the next chapter.

WHAT WOULD PROVE THIS CHAPTER WRONG. The induced-innovation loop fails if the demand-driven price surge does not convert. Concretely: if by roughly 2038 to 2040 the real cost of delivered energy per unit of useful work has not fallen below its pre-surge trend despite the current investment wave, the central mechanism of this thesis is broken and the argument should be discounted accordingly.

2. This has happened before

If the claim in chapter one were unprecedented, you would be right to dismiss it. So the first thing to establish is the precedent, and the second is to deal honestly with the strongest objection to it, which is that growth has been slowing for fifty years.

The precedent is the ordinary story of every economic leap. Two hundred years ago most of humanity worked land, and nearly everything an economy produced came from the soil. Then machines collapsed the cost of making goods, and within a few generations work moved to cities, new industries appeared that had no names before, and a kind of wealth made of intangible things came to dwarf the old kind. I want to be plain about what that transition cost, because the cheerful version of this history is a lie of omission. For roughly the first fifty years of British industrialization, output per person rose about 46 percent while real wages rose about 12: two generations of workers lived through the leap without receiving its gains, a stretch economic historians call the Engels' Pause.⁹ The gains eventually broadened, but they were broadened by institutions, unions, franchise reform, public schooling, and not delivered automatically by the machines. I will return to that in the final chapter, because it sets the honest expectation for this transition too: the pain arrives on schedule, the gains have to be fought into distribution.

Now the longer lens. Reconstructing output before the modern era is estimation work, and I will treat it as such, but the estimates agree on the shape.¹⁰ For most of recorded history, growth was so slow it was invisible within a lifetime. Then it compounded. World output growth ran at a crawl for millennia, reached roughly one percent a year only in the nineteenth century, and now adds a few percent every year. David Roodman, fitting a growth model to gross world product back to 10,000 BCE, found the long-run path is super-exponential: growth itself accelerates as output grows.¹¹ Michael Kremer's explanation is the cleanest: more people produce more ideas, ideas are non-rival, so bigger populations grow faster, and growth compounds on itself.¹² Oded Galor's unified growth theory adds the crucial shape: the acceleration is not smooth. It moves in regime shifts, long plateaus punctuated by phase changes, from the Malthusian era to take-off to modern growth.¹³ Leaps, not a ramp. Each leap in the record follows a change in the underlying inputs.

Here is the objection, and it deserves to be stated at full strength rather than survived. Since roughly 1970, growth at the technological frontier has slowed. Robert Gordon has argued that the great inventions of the second industrial revolution, electricity, sanitation, the internal combustion engine, were one-time gains that cannot recur, and that recent innovation is narrower.¹⁴ Bloom, Jones, and Van Reenen showed that ideas are getting harder to find: it takes about eighteen times more researchers to sustain Moore's law today than in the early 1970s, and research productivity has fallen across essentially every field they measured.¹⁵ The 2020s have been among the slowest growth decades since the war. Anyone selling you an acceleration story that does not engage this is selling you half a chart.

My answer is not to dispute the slowdown. It is measured, and it is real. My answer is that the slowdown and the long-run acceleration are the same story seen at two scales. Kremer's engine ran on population, and the demographic transition throttled it: the frontier's population growth stalled, and each marginal idea now costs more researcher-hours exactly as Bloom and colleagues measured. In Roodman's own model, the recent decades sit below the long-run trend line, a deviation, not a new trend.¹¹ So the question is not whether the exponent broke. The question is whether anything can restart the idea engine that demography switched off.

That is precisely the significance of AI, and it is why intelligence is the first of the four factors rather than one among equals. Aghion, Jones, and Jones worked through the growth theory: if AI automates production tasks alone, you get a step up in level; if it automates the production of ideas themselves, growth can accelerate without bound, capped only by the tasks that resist automation, the Baumol bottleneck.¹⁶ I am not resting the thesis on the explosive case. I am resting it on the narrower claim: the binding input behind the fifty-year slowdown was research effort, and research effort is the input whose cost is now collapsing.

I should state the pattern's limit honestly, because it is not true that every cost collapse produces a leap, and a version of this argument that only counts the collapses that worked is cherry-picking. Some do not cascade. The cost of nuclear fission fell, and the promised transformation never arrived, held back by accidents, regulation, and construction costs that rose faster than the physics fell. Supersonic passenger flight was achieved and then abandoned, because the cost never fell far enough for the market it needed. So a falling cost is necessary for a leap, not sufficient. What separates the collapses that cascaded, printing, steam, electricity, computing, from the ones that stalled is two things: whether the cheapened input was general, usable across the whole economy rather than one corner of it, and whether anything blocked its spread. That test is the reason this thesis rests on intelligence and energy, the two most general inputs there

are, rather than on any single technology. It is also a real risk to the forecast: if intelligence turns out to be blocked in diffusion the way fission was, the leap stalls, and the productivity condition below is where that failure would show up first.

A note on what this chapter does not establish. The long record earns the pattern, leaps happen and follow input-cost collapses, and it earns the direction. It does not earn a date. Roodman's fitted curve, extrapolated naively, reaches infinite output this century, which is impossible; the curve must bend, and whether it bends upward into a new regime now or downward into Gordon's stagnation is the live question, not a settled one. Betting against resumption has lost for ten thousand years. That is an inductive argument, and I hold it as one.

WHAT WOULD PROVE THIS CHAPTER WRONG. If, a decade into broad AI diffusion, measured productivity growth at the frontier has not risen above its 2005 to 2025 trend, then the idea-engine story fails, Gordon is right that the gains are narrow, and the leap this document forecasts is not coming on any timescale that matters to the reader.

3. The four factors, one at a time

Chapters one and two claimed that all four inputs are moving. A claim that size should be shown input by input, with each one carrying only the weight its evidence supports. The honest scorecard, before the detail: intelligence is falling now, energy is falling now while passing through an expensive phase, labour is turning, and raw material is the furthest out, a leading indicator rather than a present fact.

3.1 Intelligence

The cost collapse in machine intelligence is the fastest of the four and the best documented. The price of output at a fixed capability level has fallen between ninefold and nine-hundredfold per year depending on the task; output at the level of GPT-4, which cost tens of dollars per million tokens at launch, cost around forty cents per million by early 2026, roughly a thousandfold fall in three years.¹⁷

I want to claim exactly this and no more. Cheaper is not the same as smarter. Whether the capability frontier keeps advancing is a live scientific fight: pre-training improvements have slowed, prominent researchers argue current systems do not reason so much as pattern-match, and the rebuttals are equally serious.¹⁸ This thesis does not need to win that fight. Commoditization of the intelligence that already exists is sufficient for everything that follows, because most of the world's work does not require frontier genius. It requires competent reasoning applied cheaply, at volume, where no competent reasoning was affordable before.

The honest complication is that cheap intelligence has not yet shown up in the productivity statistics. Task-level experiments show large gains, particularly for less-experienced workers.¹⁹ Economy-wide, Acemoglu estimates AI adds only a fraction of a percentage point to productivity over a decade, an MIT survey found 95 percent of enterprise AI deployments showing no measurable return so far, and one careful field study found experienced developers slower with AI tools.²⁰ Erik Brynjolfsson's answer, which the history of electricity and computing supports, is the J-curve: general-purpose technologies depress measured productivity while firms rebuild their processes around them, then the gains arrive late and large.²¹ Both readings are open. The gap between them, three orders of magnitude in the forecasts, is the honest range, and I say so rather than picking the flattering end.

3.2 Energy

Energy has been the binding constraint on growth for most of history; every prior leap was, underneath its machinery, an energy-cost collapse. Coal for steam, oil for the assembly line. The constraint is now breaking in the open. Solar photovoltaic power with good financing is, by the International Energy Agency's assessment, the cheapest electricity in history.²² The piece that always limited solar, storage, is following: the benchmark cost of a four-hour grid battery fell 27 percent in a single year to 78 dollars per megawatt-hour, the lowest on record, and sodium-ion cells, built from one of the most abundant materials on earth, have reached rough cost parity with lithium.²³ Fusion remains a research bet, and I will not build on it; what can be said honestly is that commercial programs are now hitting physical milestones, first private deuterium-tritium operation, first direct conversion of fusion energy to electricity, rather than only publishing timelines.²⁴ I plan on solar and storage. Fusion is upside.

The bear case needs stating, because the cheap kilowatt-hour is not the whole system. Levelized cost understates the expense of firming an intermittent supply into around-the-clock power; whole-energy transitions have historically taken half-centuries, and electricity is only about a fifth of final energy use (Vaclav Smil's corrective, which I accept as a constraint on speed, not direction).²⁵ And there is Jevons: when energy gets cheaper, we use more of it, so total spending on energy may never fall even as the unit cost collapses.²⁶ I accept that too, and note that it concedes the point that matters. The claim is abundance of the input, not a smaller energy bill. An economy that uses vastly more energy because energy is cheap is the outcome this thesis predicts.

The clearest live demonstration is the country this document is written for. In 2020, solar was under three percent of Pakistan's electricity. By 2025 it was more than a fifth.²⁷ Pakistan imported more solar panels in 2024 than any country in the world, roughly seventeen gigawatts, and rooftop net-metered capacity went from under one gigawatt to over five in about two years.²⁸ No national plan drove this. Panel prices roughly halved between 2022 and 2024 while grid tariffs climbed, and millions of households and small factories did the arithmetic on their own roofs. I should be honest about the aftermath, because it illustrates the price-versus-cost rule from chapter one: in February 2026 the regulator rolled back net-metering terms, cutting what rooftop producers are paid.²⁹ The politics of the grid are now reacting against the falling cost of power. That is friction on the price. The cost keeps falling regardless.

A short word on water, because I have argued elsewhere that water reduces to energy and I want to size that claim correctly here. The largest variable cost of desalination is energy, so cheap energy does turn coastlines into water sources for cities and high-value agriculture **. But most of Pakistan's water problem is not a desalination problem. Roughly four-fifths of the country's water goes to agriculture, much of it unpriced and flood-irrigated, and the binding constraints are storage, canal losses, and groundwater depletion, none of which a cheap kilowatt-hour fixes. Energy deflation helps water at the margin. Pricing and infrastructure decide the core, and that is a chapter-seven problem.

3.3 Labour

Labour is the factor that looked immune, and it is the one I will claim most carefully. Humanoid robotics is moving from demonstration to production: one American manufacturer is converting a vehicle line to build humanoids at scale this year, and another is producing them at roughly one per hour, ramping toward thousands per year.³⁰ Against that, Moravec's paradox still holds, the machines that reason cheaply still fumble physical work a child manages, and the robots doing useful factory labour today are few. Labour is turning, on a horizon of years to a decade, not falling now, and this document treats it that way.

What is falling now is the labour value of routine cognitive work, and the evidence has a shape that matters enormously for the advice this document ends with. The displacement is landing first on entry-level work. Early payroll studies find employment for the youngest workers in the most AI-exposed occupations down on the order of thirteen percent since generative AI arrived, while older workers in the same occupations held steady.³¹ The rung juniors used to climb is the rung the machine does. I flag it here and answer it in chapter eight, because any honest version of "learn to use AI" has to face the fact that the traditional way of learning a trade, being paid to do its routine work, is the exact thing being automated.

3.4 Raw material

Raw material is the furthest factor out, and I will scope it precisely, because it is the one where enthusiasm most often outruns evidence.

The strong claim, that space mining will feed Earth's material supply, is a horizon claim and might never pencil. The first wave of asteroid-mining ventures failed commercially inside a decade, and returning bulk material to Earth may never beat terrestrial extraction on cost **. I am not building on it.

The claim I am building on is the leading indicator plus the pattern from my own industry. The cost of reaching orbit fell from tens of thousands of dollars per kilogram in the shuttle era to roughly three thousand dollars on a Falcon 9 today, with SpaceX's internal cost estimated at roughly half its customer price, which is the number that matters, per the cost-versus-price rule.³² The next vehicle generation targets a few hundred dollars per kilogram, and independent projections put roughly one hundred dollars per kilogram around 2040.³³ Those are targets and projections, not achievements, and I mark them as such.

What makes this more than a rocket story is the staircase of thresholds, because each price level switches on a layer of economy the way cheap bandwidth switched on layers of the internet. Around today's cost, communications mega-constellations and imaging-as-a-service are already profitable businesses; a Pakistani farmer can buy crop insight read from orbit today. As cost falls toward five hundred dollars per kilogram, commercial space stations, on-orbit servicing, refuelling, and early manufacturing become viable; the current station retires around 2030 and multiple private replacements are funded and under construction. Lower still, manufacturing in orbit flips from research cost to export product; a leading start-up has already manufactured pharmaceutical crystals in orbit and returned them to Earth.³⁴ Then lunar operations, and eventually tourism at prices a professional could reach.

The evidence that this staircase is real is the capital and the companies climbing it. Venture funding into space startups runs six to eight billion dollars a year, and the layer beneath the headline names now includes funded companies at every step: reusable launch (Stoke Space, half a billion dollars raised), orbital logistics (Impulse Space, over a billion), stations (Vast, Axiom, Sierra Space, collectively well over a billion), servicing and debris (Astroscale, Starfish), in-space manufacturing (Varda), and observation and data (Planet, ICEYE, Pixxel).³⁵ The detail most relevant to this document's reader: India has gone from under fifty space startups before 2020 to several hundred, produced its first launch unicorn, and its companies now span the entire value chain, launch, satellites, propulsion, surveillance, data.³⁶ A neighbouring emerging economy is participating in this frontier now, at the startup level, without owning a SpaceX. That is the model, and it is the point of this whole section: the opportunity in a deflating frontier is rarely the headline company. It is the hundreds of companies and the millions of jobs that form in the layers underneath, and those are open to people who watch the cost curve and position early.

WHAT WOULD PROVE THIS CHAPTER WRONG. Intelligence: if the real cost of a fixed unit of AI capability stops falling for three to five consecutive years. Energy: the chapter-one condition. Labour: if by the early 2030s humanoid systems remain economically marginal outside demonstrations. Raw material: if the marginal cost to orbit stalls above roughly a thousand dollars per kilogram through the 2030s, or the startup layer beneath launch stops expanding, the staircase stops, and this factor should be dropped from the thesis.

4. What this does to work

The first place falling input costs bite is the labour market, and for the reader of this document the bite is not abstract. So this chapter stays concrete: what is dying, where value moves, how to read the financial mania around it, and what happens to learning itself.

4.1 *The reckoning*

The work that dies first is output sold by the unit: the website, the logo, the batch of ads, the routine code. When a machine produces the output in seconds, no margin survives on producing it slowly. This is already visible at national scale in the world's largest offshore-services industry. Fresh-graduate hiring at India's major IT firms has collapsed, by one industry analysis from roughly six hundred thousand a year to a hundred and fifty thousand,³⁷ the largest firm announced the first mass layoff in its history,³⁸ and on the order of a hundred billion dollars came off the top firms' market value in early 2026. And yet the same firms' revenues are broadly stable or rising. That combination, revenue holding while junior headcount collapses, is the signature of what is breaking: not the industry, but the labour-arbitrage pyramid inside it, five cheap juniors sold for the price of one expensive engineer. The value is migrating up the chain, to people who design systems, direct the machines, and own outcomes. A reset, not a death, and the distinction decides everything about how to respond to it.

4.2 *From output to outcome*

When output is nearly free, the price of output goes to zero and the value moves to two things: the judgment to know which output is worth producing, and the trust and reach to deliver it as a result. Sell outcomes, not hours.

The clearest articulation comes from inside the field. Andrej Karpathy, who coined the term for AI-written software, describes the change as a floor and a ceiling: "vibe coding" raises the floor, so nearly anyone can now make software, while the engineers who direct fleets of AI systems raise the ceiling far beyond the old ten-times engineer. His summary of the human's remaining job is the thesis of this chapter: you can outsource your thinking, but you cannot outsource your understanding.³⁹ The venture firm Andreessen Horowitz called the same moment "software's YouTube moment": you no longer need to be interested in software to build software, you just need to be interested in good ideas.⁴⁰

Two examples, offered with a warning label. Cal AI, a photo-based calorie app built by two teenagers with a team of seven, passed fifteen million downloads and thirty million dollars in annual revenue in under two years.⁴¹ Kalakar, built by a South Asian team tired of captioning tools that ignored Urdu and Hindi, turned that overlooked niche into a subscription business.⁴² The warning label: these are survivors. For each of them, many equally earnest teams failed, and quoting winners as proof would be the exact survivorship error this document criticizes elsewhere. What the examples establish is narrower and still use-

ful: the cost floor is now low enough that a handful of people with domain judgment can serve a real market. The method is transferable. The outcome is not.

The scale of the reshuffle is easy to understate. Cursor, the AI coding company, went from founding to roughly a billion dollars in annualized revenue faster than any business software company in history, and in June 2026 SpaceX agreed to acquire it for sixty billion dollars, a deal pending close as I write.⁴³ One of its four co-founders is a Karachi-born mathematician in his twenties.⁴⁴ I cite the revenue as the real signal and the acquisition price as something else, which brings me to the bubble.

4.3 Reading the bubble

Two things are true at once, and holding them apart is the most practical financial skill of this decade. The technology is real. The financing around it is a mania that will correct.

The mania is not subtle. OpenAI has committed roughly 1.4 trillion dollars to data-centre buildouts against roughly twenty billion in annualized revenue.⁴⁵ The money moves in a circle, chip-maker investing in model-maker who buys the chips, that analysts compare to the vendor financing of the dot-com era.⁴⁶ Five companies make up about a third of the S&P 500, the highest concentration in half a century.⁴⁷ An MIT survey found 95 percent of enterprise AI pilots with no measurable return yet.⁴⁸ People inside the industry say plainly that someone will lose a phenomenal amount of money. They are right.

And none of it is evidence against the technology, because this is what transformative technologies look like in their bubble phase, every time. British railway shares fell by two-thirds in the 1840s; the railways reshaped the world for a century afterward. The dot-com crash erased five trillion dollars and seventy-eight percent of the Nasdaq; roughly half the companies survived, and the real ones, including one that fell about ninety percent peak to trough, became the most valuable firms in history.⁴⁹ The telecom bust bankrupted the operators and left behind the fibre the internet runs on. The lesson is mechanical, not motivational: the crash reprises the companies, not the input costs. Cheap intelligence, like cheap fibre, survives its own bubble. So the strategy writes itself. Build for the revolution, survive the bubble, and when the narrative flips overnight from mania to "AI was a fraud," change nothing, because neither narrative was about the cost curves.

4.4 Education returns to mastery

One more consequence of cheap intelligence, and the one I find most personally important. Mass schooling in its current form, one lesson delivered to forty children at once, is an industrial-era design; before it, the few who were educated at all were tutored one-to-one, to mastery.⁵⁰ In 1984 Benjamin Bloom measured what that older model is worth: students tutored one-to-one with mastery methods performed around two standard deviations above conventional classrooms, a result later replications temper to something like half a standard deviation, large still, with the direction never in dispute.⁵¹ The barrier was always cost: no society could afford a tutor per child. Cost is the one variable this whole document is about. AI is the first technology that makes one-to-one mastery tutoring a mass-market product rather than an aristocrat's privilege, and for a country where more than twenty-five million children are out of school and public education spending has fallen to under one percent of GDP,⁵² that is not an edtech story. It is the difference between a generation that participates in what follows and one that watches it.

WHAT WOULD PROVE THIS CHAPTER WRONG. If, five years from now, the offshore-services reset has reversed (junior hiring recovered to trend without role redesign), if outcome-priced work has not grown its share

against hourly work, or if AI tutoring at scale shows no measured learning gains over conventional instruction, then value is not moving the way this chapter claims.

5. How economic power moves

Individuals and firms respond to input costs quickly. Countries respond slowly, and whole eras respond slowest of all. This chapter zooms out to the largest scale, and I will keep it deliberately cool, because this ground is usually argued with either triumph or grievance and both make people stupid.

Start with the plain claim, because it is simpler than the scholarship around it. No country stays on top forever, and which one is on top tends to track who controls the era's most important and cheapest inputs. When those inputs change, the ranking eventually changes with them. That is not a fringe idea or a motivational one. It is one of the oldest structural claims in historical scholarship, argued in different forms by Braudel, Wallerstein, Arrighi, and Kennedy, and more recently and more controversially by Turchin and Ray Dalio.⁵³ They disagree with each other about the mechanism, the timing, and whether the handovers are peaceful, and grand cyclical history has an embarrassing record whenever it names a date, Spengler being the standing warning. So I will take from it only what it can bear: dominance has never been permanent, its shifts track underlying production and cost, and treating the current arrangement as the final one is the single position with a perfect record of being wrong.

The long data adds one fact most readers of English were never taught. By the best available estimates, the Indian subcontinent and China together produced roughly half the world's output through the sixteenth and seventeenth centuries; the balance tipped decisively toward the Atlantic world only around 1800.⁵⁴ Two honest cautions. Pre-modern GDP is reconstruction, with wide error bands. And the old Eastern weight was size, population multiplied by subsistence, not per-capita wealth: by 1600 an average Briton already out-earned an average Indian, and the gap widened for three centuries afterward.⁵⁵ So the correct claim is not "the East was rich and will be rich again." It is that the current distribution of economic weight is younger than it looks, and the mechanisms that produced it, cheap energy captured first, industrial scale, control of distribution, are exactly the mechanisms now being repriced. The measured centre of gravity of world economic activity has in fact been moving east for forty years, from the mid-Atlantic in 1980 to Central Asia's longitude on current projections.⁵⁶

The strains on the incumbent order are real and specific, and they should be stated without relish. United States gross interest payments on federal debt passed a trillion dollars a year in fiscal 2025, with roughly a quarter flowing to foreign creditors.⁵⁷ Inequality has produced what its own analysts call a K-shaped economy. But disaggregate before concluding: the demographic decline that constrains Europe, Japan, and China applies far less to the United States, which projects a growing workforce for decades, and the American innovation and energy base remains the deepest in the world. Decline-of-the-West is a much weaker claim than rebalancing-of-weight, and this document makes only the second.

Within that frame, current American strategy is legible rather than mysterious. The stated policy has turned explicitly toward hemispheric consolidation, the 2025 National Security Strategy invokes the Monroe Doctrine by name,⁵⁸ and a serious school of strategists has argued for exactly this retrenchment for years. Read through chapter one's economics, the deeper logic is factor prices: the developed world's binding costs were labour and intelligence, which it imported for decades. AI and robotics are the first technologies that promise to produce both domestically. Robert Allen's explanation of why Britain indus-

trialized first, dear labour plus cheap energy makes automation profitable, is operating again in reverse geography.⁵⁹ Whether it works is an open question: the robot density, skills base, and supplier ecosystems are not yet there, manufacturing employment fell in 2025 rather than rose, and I treat reindustrialization as a fifteen-to-twenty-year bet, not a five-year fact.⁶⁰ But the direction of effort is public policy, not conjecture. And its corollary matters more to my reader than the policy itself: a great power turning toward its own hemisphere loosens the patronage structures it maintained elsewhere. That is chapter six.

WHAT WOULD PROVE THIS CHAPTER WRONG. If automated domestic manufacturing in high-wage economies is not cost-competitive with offshore production within fifteen to twenty years, the reindustrialization leg fails. If the eastward drift of economic activity's centre of gravity stalls or reverses over the coming decade, the re-balancing leg fails, and with it much of chapter six.

6. The opening, and why it has to be earned

Everything to this point has been about the world. This chapter is about the reader's part of it, and it is the chapter where I am most determined not to sell you anything, because the case for hope here is strong enough without inflation, and the case against complacency is stronger still.

Start with the despair, because it is rational at the scale it is usually felt. If your reference frame is the last few decades, Pakistan's story is falling behind: crises, bailouts, emigration queues. A fifth of the world's out-of-school children. The despair reads the recent past correctly. What it misreads is the arc, in both directions: backward, where the region's poverty is an anomaly of the last two centuries rather than its permanent condition, and forward, where the repricing described in this document is precisely the kind of discontinuity in which relative positions change. Goldman Sachs's long-horizon projections, for what such exercises are worth, put Pakistan among the world's six largest economies by 2075, and the authors' own caveat is the honest one: the projection is demographic arithmetic, conditional on institutions and policy delivering, which is exactly the condition in question.⁶¹

So the real question is institutional, and I will state my view of it carefully, as a mechanism rather than a prophecy.

High-inequality, rent-extracting political orders are stable so long as two supports hold: a coercive apparatus capable of containing discontent, and, frequently, an external patron with reasons to fund that apparatus.⁶² The political-science literature on when such orders break is large and sobering. Revolutions ignite not at maximum misery but when rising expectations meet reversal, Davies's J-curve; large young cohorts raise the statistical risk of instability substantially, and Pakistan has one of the largest youth bulges on earth; elite overproduction and popular immiseration move together ahead of state crises; and, the finding most relevant here, elites historically concede reform precisely when the threat of rupture becomes credible.⁶³ Meanwhile the strains on the old supports are visible: the patron's attention is moving hemispheres, the population is younger, more connected, and harder to govern than any previous generation, and the recent regional record, Sri Lanka in 2022, Bangladesh in 2024, Nepal after it, shows what connected youth cohorts now do to governments that fail them.

What I will not do is call the outcome. The same literature is unambiguous that rupture is not reform: of the five main Arab Spring states, four ended in civil war or harder authoritarianism, and Walter Scheidel's grim survey found that historically, large levellings of inequality came through catastrophe far

more often than through orderly renegotiation.⁶⁴ So the claim I am entitled to is this: the current equilibrium is becoming more expensive to hold, pressure toward renegotiation is accumulating, and the window for a better social contract will open, likely more than once. Whether the country passes through it is not written anywhere. Windows are not outcomes.

The same discipline applies to the economics. The base rate against catch-up growth is brutal: developing economies have diverged from rich ones more often than converged, growth slowdowns cluster exactly in the middle-income range, and only thirteen economies since 1950 have sustained the growth rates that transform a country within a generation.⁶⁵ The demographic dividend that optimists cite is real but conditional, East Asia converted its youth bulge because the schools, jobs, and institutions were ready, and Pakistan currently fails those preconditions: the out-of-school millions, education spending under one percent of GDP, a quarter of women in the workforce.⁵² Anyone who tells you Pakistan's rise is inevitable is doing to you what the crypto sellers did.

I will make one part of this concrete, because I have a personal stake in it and it is the kind of thing that is easy to miss under the political noise. I have two young children, and like any parent I worry about what is in their food, the pesticides and preservatives you cannot wash off. Organic food at the scale a country needs has always been a contradiction, too expensive to feed everyone. The same cheap intelligence this paper is about is already being pointed at that problem: small teams building targeted systems, lasers that kill insects without spraying, sensors that let a farmer grow more with less, that turn organic-quality food at commercial scale from an impossibility into a cost problem, which is the kind of problem that gets solved. While the argument on television is about who to blame for the last decade, people are quietly building the next one. Most of this paper is a case for noticing them, and then for being one of them.

Here is what survives that honesty, and it is not small. First, the entry ticket has collapsed. Catch-up used to require what only states could build, ports, factories, capital stock. The current repricing runs on inputs an individual can buy: intelligence at pennies, distribution for free, energy off your own roof. Pakistan's IT exports are growing twenty percent a year and its freelancers, over two million of them, earned a record 856 million dollars in nine months of the current fiscal year, half again more than the year before, under a registered-exporter tax of a quarter of one percent.⁶⁶ These are small numbers against the need and I will not pretend otherwise; the point is their slope, and that they are being produced bottom-up, by individuals, ahead of any institutional reform, exactly like the rooftop solar. Second, the absence of legacy is, for once, an advantage. India's IT giants must defend a pyramid that AI is dissolving; a nineteen-year-old in Gujranwala has nothing to unlearn. Third, the precedent next door is instructive in both directions: India in 1991 was pledging its gold reserves to avoid default, chose the unglamorous digital plumbing work, and compounded it into services exports approaching the size of Pakistan's entire economy.⁶⁷ The path exists. It was walked from a worse starting point than Pakistan's. And it took thirty years of compounding, which is the honest price tag.

WHAT WOULD PROVE THIS CHAPTER WRONG. Concretely: if by 2035 Pakistan's IT and digital-services exports have not at least tripled from today's roughly four billion dollars, or if freelancer earnings have flattened rather than continued their recent double-digit annual growth, the individual-level opening was overstated. And if, over the same decade, the region's pressure toward renegotiation produces only harder entrenchment with no episode of institutional opening in any comparable state, then the political mechanism was misread and the despair was the better guide.

7. Money is the language of value

One belief runs underneath the previous chapter's failures, the unpriced water, the unreformed institutions, the wasted dividend, and it is cultural before it is economic. It deserves its own short chapter because I hold it to be Pakistan's deepest economic problem and because it is the belief this document most directly challenges.

The belief is that money is morally suspect: that to price something is to cheapen it, that the money-minded are the corrupted, that the good things, water, education, care, are degraded by being valued in rupees. The correction is one sentence long. Money is not the opposite of value. It is the language value is spoken in, the instrument a society uses to make value legible so it can be allocated, conserved, and grown. This is the settled view of value in economics since the marginal revolution, price as society's running measurement of what things are worth relative to each other,⁶⁸ and its deepest philosophical treatment is Georg Simmel's: money as the medium that lets incommensurable things be compared and exchanged at all.⁶⁹

A society that refuses to speak the language does not thereby protect its values. It loses the ability to manage them. The master case is the one from chapter three: Pakistan calls water precious and prices it at nearly nothing, so four-fifths of it floods fields growing water-hungry crops, storage goes unbuilt, and foreign agribusiness harvests the un-priced subsidy and ships it abroad as crops. The refusal to price the resource is the mechanism of its waste. Hardin called the general case the tragedy of the commons; Pigou supplied the fix a century ago, price the thing at its true scarcity.⁷⁰

Now the two objections, because this idea is only safe to hold with both hands. Elinor Ostrom won the Nobel prize for showing that commons can be governed well without pricing or privatization, by communities with clear boundaries and rules, and many of her cases are irrigation systems.⁷¹ Michael Sandel adds that some goods are corrupted by pricing, the Israeli day-care study where fining lateness increased it, because the fine converted a moral duty into a cheap fee.⁷² Both are right, and they scope the claim rather than defeat it: the requirement is that value be made legible and consequential, by price or by binding rules, at the scale where anonymity defeats social enforcement, national water allocation, energy tariffs, land, the price mechanism is usually the only language strong enough. What is never acceptable is the current arrangement, which is neither: value declared sacred in speech and treated as free in practice.

One more connection, and the chapter's real point. The cost-versus-price rule from chapter one is this same idea seen from the other side. Where entry is open, price speaks cost honestly. Where a cartel, a monopoly, or a patronage structure blocks entry, price stops speaking, the falling cost accumulates as extracted margin, and society experiences abundance as continued scarcity. A country that moralizes against money while tolerating unpriced water and rigged prices has the worst of both worlds: it neither speaks the language nor polices the liars. Learning to speak it, plainly, without embarrassment, is not a surrender to greed. It is how the falling costs described in this document reach the people they are falling for.

WHAT WOULD PROVE THIS CHAPTER WRONG. Concretely: if over the next decade the states that have moved to price water at national scale, Australia, Israel, parts of India, show measurably worse conservation and allocation outcomes than comparable systems that kept it unpriced, then price is not the tool I claim it is at that scale, and Ostrom's exception is closer to the rule.

8. What to do about it

Everything in this document narrows to a decision that is personal before it is national, so I will end at the scale of one person, most likely you.

Why intelligence is the first move. Of the four factors, intelligence is the one falling fastest, the one that accelerates the other three, and the only one you can buy tonight for the price of a meal. Energy needs capital, robotics needs years, orbit needs a rocket. Cheap reasoning is already in your pocket. That is why the current moment, before the diffusion completes and the advantage normalizes, is the widest part of the window.

The line, and its uncomfortable half. AI will not take your job; the person using AI will. I have said it for years and the evidence has kept it true. But chapter three surfaced the uncomfortable half: the entry-level rung, the paid apprenticeship of routine work, is the part being automated first.³¹ The old sequence, junior job, then skill, then judgment, is breaking at its first step. I refuse to soften this, because it is precisely why moving early matters: the people who build judgment now, while the old rungs still partially exist, will be the ones who never needed them. The ones who wait will face the ladder with the bottom missing. The cruelty of that timing is real, it is not your fault, and it changes nothing about what you should do.

The corollary I have lived. When building gets cheap, value moves to what cannot be commoditized: knowing what is worth building, and being trusted by the people it is for. Taste, narrative, distribution. I watched this happen in media, nobody who mattered cared about your camera by 2015, they cared whether you had something to say and an audience that believed you, and it is now happening to software and, in sequence, to every domain the four factors touch. Building your judgment and your audience is not marketing advice. It is positioning for where value provably migrates when production commoditizes.

The method. It has four steps and no shortcuts. Find a real problem you or your community has; your domain knowledge, the industry you know, the city you live in, is the unfair advantage that makes the problem visible to you and invisible to Silicon Valley. Build the smallest solution that works, with AI doing the labour that used to require a team. Show the work publicly, so the people with the problem can find you and the trust can compound. Charge, because revenue is the only feedback that cannot flatter you. Then repeat, and let each pass raise your judgment. You do not need a million users; you need one real problem and a recurring reason to be paid. The examples in chapter four were survivors, and I promised not to sell you survivorship: most attempts fail, mine failed repeatedly, and the method's honest promise is not success but compounding, each iteration leaves you with more skill, more audience, and more judgment than the last, which is more than any failed job application leaves behind.

The temperament. Everything above is publicly known, and still most people will not do it, because the returns compound quietly and the noise pays instantly. The discipline that matters is not intensity but duration: the willingness to be two years into something that looks like nothing from outside. My audience grew invisibly for two years before anything happened. The people who survived media were not the most talented; they were the ones still there when the compounding arrived. Patience over urgency, the long term over the visible term, building value over extracting it. I hold these as practical instructions, not virtues.

The steelman, last. The strongest case against this document, stated better than its critics usually state it: the whole argument could be survivorship bias wearing a framework, one lucky creator generalizing his

luck; reasoning by historical analogy proves nothing about the future; the growth data since 1970 points down, not up; revolutions mostly end badly; catch-up growth mostly fails; and the document's optimism conveniently flatters the audience it is written for. I accept every one of those as the correct prior. The document's answer is not that the pessimistic priors are wrong; it is that each has a specific, named condition under which it flips, falling input costs, a restarted idea engine, a collapsed entry ticket, and those conditions are now measurably in motion, with kill-dates attached at the end of every chapter of this document. If they fail, I have told you exactly where to watch it fail. That is the difference I can offer between this and hope: hope does not tell you what would falsify it.

There is one objection I am the worst-placed person in the world to answer, and I would rather name it plainly than let it sit unspoken. I make my living from the exact posture this document recommends: build early, own your distribution. That is a conflict of interest, and you should discount what I say accordingly. I also cannot honestly rule out that I am one lucky person generalizing his own luck into a law, because I am inside the sample I am reasoning from. So let me separate the two things I am claiming, because they can fail independently. The first is an analysis: the cost of these inputs is falling, and that has dated, quantified tests at the end of every chapter, none of which depend on me being right about what to do. The second is advice: that the response is to participate early. That one has its own clean test, and it cuts against my interest to state it. If, over the next decade, the input costs fall as I claim but the people who moved early do no better on average than those who waited, then the advice was wrong even though the analysis was right, and this document's conclusion should be discarded while its diagnosis stands. I cannot prove I am not the lucky exception. I can tell you the number that would prove my advice was survivorship after all.

And if the conditions hold, then the arithmetic of this moment is simple. The cost of thinking is collapsing. The cost of power is collapsing. The doors that were held shut by the price of entry are opening in sequence, and they are opening widest exactly where they were locked longest. Nothing about that guarantees you anything. It only guarantees that, for the first time in generations, the limiting factor is no longer the price of the tools.

It is what you do with them.

References

Draft convention: books and journal articles are cited in standard form and need no URL. Web-sourced figures carry the URLs verified in the research pass ('dossier/00-fact-verification-layer.md' and section files). Items marked have their exact source in the named dossier file and are resolved in the final citation pass.

Additional working sources for every chapter, including the contrarian literature engaged above, are catalogued in the research dossier ('docs/reference/dossier/'), which this document was built from.

1. Justus von Liebig's "law of the minimum" (agronomy, 1840s), as the general scarcest-input principle; Michael Kremer, "The O-Ring Theory of Economic Development," *Quarterly Journal of Economics* 108(3), 1993.
2. John R. Hicks, *The Theory of Wages* (Macmillan, 1932), the induced-innovation hypothesis.
3. Daron Acemoglu, "Directed Technical Change," *Review of Economic Studies* 69(4), 2002.
4. David Popp, "Induced Innovation and Energy Prices," *American Economic Review* 92(1), 2002.
5. Richard G. Newell, Adam B. Jaffe, Robert N. Stavins, "The Induced Innovation Hypothesis and Energy-Saving Technological Change," *Quarterly Journal of Economics* 114(3), 1999.

6. US electricity prices +6.9% in 2025 (~2x inflation); data centres ~40% of demand growth; >\$29B in utility rate cases H1 2025.
7. IEA, *World Energy Investment 2025*: record ~\$3.3T total, ~\$2.2T clean energy, acceleration attributed substantially to AI/data-centre demand.
8. Three Mile Island restart PPA; Google agreements with Kairos Power (SMR) and Commonwealth Fusion Systems.
9. Robert C. Allen, "Engels' Pause: Technical Change, Capital Accumulation, and Inequality in the British Industrial Revolution," *Explorations in Economic History* 46(4), 2009.
10. Maddison Project Database (Bolt & van Zanden, 2023 update), <https://onlinelibrary.wiley.com/doi/full/10.1111/joes.12618>
11. David Roodman, "Modeling the Human Trajectory," Open Philanthropy, 2020, <https://www.openphilanthropy.org/research/modeling-the-human-trajectory/>
12. Michael Kremer, "Population Growth and Technological Change: One Million Years to 1990," *Quarterly Journal of Economics* 108(3), 1993.
13. Oded Galor, *Unified Growth Theory* (Princeton University Press, 2011).
14. Robert J. Gordon, *The Rise and Fall of American Growth* (Princeton University Press, 2016).
15. Nicholas Bloom, Charles I. Jones, John Van Reenen, Michael Webb, "Are Ideas Getting Harder to Find?" *American Economic Review* 110(4), 2020.
16. Philippe Aghion, Benjamin F. Jones, Charles I. Jones, "Artificial Intelligence and Economic Growth," NBER Working Paper 23928 (2017); in *The Economics of Artificial Intelligence* (2019).
17. Epoch AI, "LLM inference prices have fallen rapidly but unequally across tasks," <https://epoch.ai/data-insights/llm-inference-price-trends> ; corroborating cost analyses (a16z "LLMflation").
18. The capability-plateau debate: Apple, "The Illusion of Thinking" (2025) and the Lawszen rebuttal; Gary Marcus's critiques.
19. Task-level studies: Brynjolfsson, Li & Raymond (call-centre agents, ~14% productivity, largest for novices); Noy & Zhang (writing tasks).
20. Daron Acemoglu, "The Simple Macroeconomics of AI" (2024), ~0.66% TFP over 10 years; MIT Project NANDA, "The GenAI Divide: State of AI in Business 2025" (95% zero measurable return), via <https://fortune.com/2025/08/18/mit-report-95-percent-generative-ai-pilots-at-companies-failing-cfo/> ; METR developer study (19% slower).
21. Erik Brynjolfsson, Daniel Rock, Chad Syverson, "The Productivity J-Curve," *American Economic Journal: Macroeconomics* 13(1), 2021.
22. IEA, *World Energy Outlook 2020* ("cheapest source of electricity in history"), via <https://www.carbonbrief.org/solar-is-now-cheapest-electricity-in-history-confirms-iea/>
23. BloombergNEF, LCOE 2026: 4-hour battery benchmark \$78/MWh, -27% YoY, record low, <https://about.bnef.com/insights/clean-energy/battery-storage-costs-hit-record-lows-as-costs-of-other-clean-power-technologies-increased-bloombergnef/> ; sodium-ion near-parity, <https://www.pv-magazine.com/2026/01/09/sodium-ion-battery-cells-already-near-lithium-ion-cost-parity-set-to-get-cheaper/>
24. NIF ignition (DOE, 2022 onward), <https://www.energy.gov/articles/doe-national-laboratory-makes-history-achieving-fusion-ignition> ; Realta Fusion direct energy conversion (June 2026), <https://www.ans.org/news/2026-07-06/article-8179/> ; Helion first private D-T operation.
25. Vaclav Smil, *Energy Transitions: Global and National Perspectives* (2017); system-cost critiques (Idel, LFSCOE; Hirth, integration costs).
26. W. Stanley Jevons, *The Coal Question* (1865); the modern rebound literature.
27. Ember, "The solarisation of Pakistan's energy economy" (solar ~22% of generation H1 2025, from ~3% in 2020), <https://ember-energy.org/latest-insights/the-solarisation-of-pakistans-energy-economy/>
28. Pakistan world's largest panel importer 2024 (~17 GW, Ember Global Electricity Review 2025); rooftop net-metered <1 GW (2023) → ~5.3 GW (2025), <https://www.arabnews.com/node/2622248/amp> ; panel prices roughly halved 2022-2024, <https://propakistan.pk/2024/05/10/>
29. NEPRA Prosumer Regulations, February 2026 (net-metering rollback), <https://profit.pakistantoday.com.pk/2026/02/10/>
30. Tesla Fremont Optimus line conversion (production from ~late July/August 2026), <https://electrek.co/2026/04/22/tesla-optimus-production-fremont-model-sx-line/> ; Figure BotQ at ~1 robot/hour, <https://www.figure.ai/news/botq>

31. Erik Brynjolfsson et al., "Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of AI" (Stanford, 2025): employment of 22-25-year-olds in the most AI-exposed occupations down ~13%.
32. Falcon 9 ~\$3,000/kg customer price (list \$69.75-74M / ~22.8t LEO), internal/marginal cost estimated near half of that, <https://spacenexus.us/guide/space-launch-cost-comparison>
33. Starship sub-\$500/kg as design target; Citigroup space report, ~\$100/kg base case by 2040.
34. Varda Space Industries, first commercial in-space pharmaceutical manufacturing with successful re-entry (2024); \$187M Series C (2025).
35. Startup funding landscape 2024-2026: VC ~\$6-8B/yr into space tech; Stoke Space ~\$480M; Impulse Space >\$1B across Series C/D; Vast \$500M (2026); Axiom ~\$350M; Sierra Space ~\$290M; Astroscale, Starfish, Planet, ICEYE, Pixxel (NRO contract, Google investment). Sources: company and press announcements compiled in the author's research document; individual URLs resolved in the citation pass.
36. India: <50 spacetechn startups pre-2020 → 250-400 by 2025-26; funding ~\$196M in 2025 (~2x YoY); Skyroot \$1.1B valuation (first Indian space unicorn, May 2026).
37. GetGenerative.ai industry analysis: Indian IT fresher hiring ~600,000 (FY2022-23) → ~150,000, <https://www.getgenerative.ai/india-it-pyramid-collapse-ai/> (single-outlet figure; direction corroborated by top-4-firm hiring data).
38. TCS ~12,000 layoffs announced 2025, first mass layoff in company history; sector revenues broadly stable ("a reset, not a death"), https://www.business-standard.com/industry/news/tcs-infosys-lay-off-indian-it-sector-job-cuts-automation-ai-reskilling-125072800505_1.html
39. Andrej Karpathy, Sequoia AI Ascent 2026 write-up (floor/ceiling; "you can outsource your thinking, but you can't outsource your understanding"), <https://karpathy.bearblog.dev/sequoia-ascent-2026/>
40. Anish Acharya, "Software's YouTube Moment Is Happening Now," Andreessen Horowitz, <https://a16z.com/software-youtube-moment-is-happening-now/>
41. TechCrunch, "MyFitnessPal has acquired Cal AI..." (15M+ downloads, \$30M+ ARR, team of seven), <https://techcrunch.com/2026/03/02/myfitnesspal-has-acquired-cal-ai-the-viral-calorie-app-built-by-teens/>
42. Kalakar (South Asian AI captioning; company-stated traction), <https://www.kalakar.io/>
43. Cursor/Anysphere: fastest to ~\$1B ARR in business-software history; SpaceX agreed to acquire for \$60B, June 16 2026, all-stock, pending Q3 close, <https://www.cnbc.com/2026/06/16/spacex-spex-cursor-acquisition-ipo.html>
44. Forbes profile, Sualeh Asif (Karachi-born co-founder, IMO 2016-18, MIT), <https://www.forbes.com/profile/sualeh-asif/>
45. Sam Altman: ~\$1.4T in data-centre commitments against ~\$20B annualized revenue (Nov 2025), <https://techcrunch.com/2025/11/06/sam-altman-says-openai-has-20b-arr-and-about-1-4-trillion-in-data-center-commitments/>
46. Nvidia-OpenAI circular financing (up to \$100B investment; OpenAI purchases Nvidia systems), <https://openai.com/index/openai-nvidia-systems-partnership/>; dot-com vendor-financing comparisons (Bloomberg et al.).
47. Five largest firms ≈ one-third of S&P 500 market cap, highest concentration in ~50 years, <https://www.visualcapitalist.com/charted-sp-500-market-concentration-over-145-years/>
48. MIT Project NANDA (2025), see ref. 20.
49. Railway Mania (index fall >65%): https://en.wikipedia.org/wiki/Railway_Mania; dot-com record (Nasdaq -78%, ~\$5T erased; ~48% of dot-coms survived to 2004, Goldfarb & Kirsch, U. Maryland; Amazon ~90% peak-to-trough), https://en.wikipedia.org/wiki/Dot-com_bubble
50. On the industrial origins of mass schooling, stated in its defensible form (standardized, age-graded delivery), with the historiographic caution of Audrey Watters, "The Invented History of 'The Factory Model of Education'" (2015), <https://medium.com/the-history-of-the-future-of-education/the-invented-history-of-the-factory-model-of-education-a069ae3d1e99>
51. Benjamin S. Bloom, "The 2 Sigma Problem," *Educational Researcher* 13(6), 1984, <https://journals.sagepub.com/doi/10.3102/0013189X013006004>; replication reviews placing mastery-learning effects nearer ~0.5σ, <https://nintil.com/bloom-sigma/>
52. UNICEF Pakistan (25.1M children aged 5-16 out of school; education spending ~0.8% of GDP, record low), <https://www.unicef.org/pakistan/education>; <https://propakistani.pk/2026/06/11/>
53. Fernand Braudel, *Civilization and Capitalism* (1979); Immanuel Wallerstein, *The Modern World-System* (1974-); Giovanni Arrighi, *The Long Twentieth Century* (1994) and *Adam Smith in Beijing* (2007); Paul Kennedy, *The Rise and Fall of the*

Great Powers (1987); Peter Turchin, *Ages of Discord* (2016); Ray Dalio, *Principles for Dealing with the Changing World Order* (2021).

54. Maddison-tradition estimates: India + China $\approx$ half of world GDP c. 1600-1700, [https://en.wikipedia.org/wiki/List_of_regions_by_past_GDP_\(PPP\)](https://en.wikipedia.org/wiki/List_of_regions_by_past_GDP_(PPP))
55. Stephen Broadberry & Bishnupriya Gupta on early-modern Anglo-Indian per-capita divergence (Indian GDP per capita from >60% of Britain's c.1600 to <15% by 1871).
56. Danny Quah, "The Global Economy's Shifting Centre of Gravity," *Global Policy* 2(1), 2011; Richard Baldwin, *The Great Convergence* (2016).
57. CRFB, "Trillion-Dollar Interest Payments Are the New Norm" (gross interest >\$1T FY2025; net interest \$970B), <https://www.crfb.org/blogs/trillion-dollar-interest-payments-are-new-norm> ; CRS on foreign holdings (~\$282B interest to foreign creditors), <https://www.congress.gov/crs-product/RS22331>
58. United States National Security Strategy, 2025 (explicit Monroe Doctrine framing).
59. Robert C. Allen, *The British Industrial Revolution in Global Perspective* (Cambridge University Press, 2009).
60. Reshoring realism: 2025 US manufacturing employment decline; US robot density ~295-307 per 10,000 vs China ~470 (IFR).
61. Kevin Daly & Tadas Gedminas, "The Path to 2075," Goldman Sachs Global Economics Paper (Dec 2022): Pakistan projected 6th-largest economy by 2075, conditional on policy/institutions, <https://www.goldmansachs.com/insights/articles/how-india-services-economy-became-a-world-leader> (Goldman insights hub) [exact paper pin: dossier/01]
62. Eva Bellin, "The Robustness of Authoritarianism in the Middle East," *Comparative Politics* 36(2), 2004, and "Reconsidering the Robustness of Authoritarianism," *Comparative Politics* 44(2), 2012.
63. James C. Davies, "Toward a Theory of Revolution," *American Sociological Review* 27(1), 1962; Ted Robert Gurr, *Why Men Rebel* (1970); Jack A. Goldstone, *Revolution and Rebellion in the Early Modern World* (1991); Henrik Urdal, "A Clash of Generations? Youth Bulges and Political Violence," *International Studies Quarterly* 50(3), 2006; Daron Acemoglu & James A. Robinson, *Economic Origins of Dictatorship and Democracy* (2006).
64. Walter Scheidel, *The Great Leveler* (Princeton University Press, 2017); Arab Spring outcome record compiled in dossier/03 addendum.
65. Lant Pritchett, "Divergence, Big Time," *Journal of Economic Perspectives* 11(3), 1997; Barry Eichengreen, Donghyun Park, Kwanho Shin, "Growth Slowdowns Redux" (NBER, 2013); Commission on Growth and Development, *The Growth Report* (2008): 13 economies sustained 7%+ for 25+ years.
66. Pakistan IT exports \$3.8B FY2025 (+18%, record; third-largest FX earner), <https://www.dawn.com/news/1925145> ; 9M FY2026 ~\$3.39B (+20%), <https://bloompakistan.com/pakistan-it-exports-hit-record-3-4-billion-in-nine-months-of-fy2025-26/> ; freelancers \$856M in 9M FY26 (+50%; ~2.37M registered, ADB estimate), <https://profit.pakistantoday.com.pk/2026/04/21/pakistan-freelancers-earnings-rise-50-to-856-million-in-9mfy26/> ; 0.25% final tax for PSEB-registered IT exporters (1% unregistered), <https://pkrevenue.com/final-tax-regime-continues-for-service-exports-in-pakistan/>
67. India 1991 crisis (reserves $\approx$ 3 weeks of imports; 67 tonnes of gold pledged), https://en.wikipedia.org/wiki/1991_Indian_economic_crisis ; Goldman Sachs, India services exports \$340B (2023) $\rightarrow$ ~\$800B (2030), <https://www.ibef.org/news/india-global-services-export-to-reach-us-800-billion-by-2030-goldman-sachs>
68. Carl Menger, *Principles of Economics* (1871), and the marginal revolution (Jevons, Walras).
69. Georg Simmel, *The Philosophy of Money* (1900).
70. Garrett Hardin, "The Tragedy of the Commons," *Science* 162, 1968; A.C. Pigou, *The Economics of Welfare* (1920).
71. Elinor Ostrom, *Governing the Commons* (Cambridge University Press, 1990).
72. Michael J. Sandel, *What Money Can't Buy* (2012); Uri Gneezy & Aldo Rustichini, "A Fine Is a Price," *Journal of Legal Studies* 29(1), 2000.

About the author

Muzamil Hasan is the founder of arc. and the host of the podcasts *Thought Behind Things* and *Beyond the Bubble*. Over the past fifteen years he has built a large audience from Pakistan, helped more than five hundred people build their personal brands, and now builds the tools and the community to run the method described in this paper at scale. The argument here is the bet he has placed his own work on.

This is a working paper, meant to be read, argued with, forwarded, and improved. If it reached you from someone else, and you want to follow the work, disagree with it, or build on it, the door is open. Web: muzamilhasan.com. Email: hello@muzamilhasan.com. Podcasts: *Thought Behind Things* and *Beyond the Bubble*.

Figures were verified against primary and reputable secondary sources as of mid-2026; where a claim rests on a heuristic, an estimate, or a contested interpretation, the text says so plainly. This is a thinking document, not investment, legal, or career advice. The decisions, and the seat at the table, are yours to take.