Growth Academy

The University of Chicago
World Bank Group

Growth Academy 2026

Lecture Notes

University of Chicago & World Bank Group

Chicago • July 27–31, 2026
Washington, D.C. • August 3–5, 2026

Summaries distilled from the lectures of the 2026 Growth Academy Summer School.

Contents

Day 1 • Monday, July 27, 2026 • Chicago

The Middle-Income Trap

A lecture by **Somik Lall**, Co-Director, Growth Academy; Director, Strategy and WBG Institute for Economic Development, Development Economics Vice Presidency, World Bank Group.

The bottom line. Middle-income economies rarely stall for lack of ideas; they stall because incumbents preserve the status quo, resources sit trapped in unproductive firms, and countries never switch growth models. Escaping the trap means moving in sequence from investment to infusion to innovation, and rebalancing Schumpeter's forces — by disciplining incumbents, rewarding merit, and capitalizing on crisis.

Two traps, not one

World Bank Group's 2024 World Development Report was titled *The Middle-Income Trap*, but the sharper frame is creative destruction. The term was coined in 2006 by World Bank economists in Latin America — ground zero for economies that surge toward the frontier and then slow. That slowing is measurable: a growth slowdown is about three times more likely in a middle-income economy than in an advanced one, and slowdowns arrive sooner and last longer where economic freedom is lower.

The pattern is stubborn. Since 1980, middle-income countries' average per-capita income relative to the United States has barely moved, near six percent, with China the only real bright spot. These economies hold about three-quarters of the world's population but only about a third of its output, yet nearly all aspire to join the club of developed nations, not merely to clear the Bank's high-income threshold of about $14,000 — a low bar, since the United States sits roughly seven times higher. The transition is hard because there was never a real theory of growth for the middle: Solow–Swan models fit low-income economies and Lucas–Romer models fit advanced ones, but neither explains the switch — and the trap is really two, one near $4,000–$5,000 and another near $8,000–$10,000, where countries stall for failing to change models.

The real bottleneck: misallocation of capital, talent, and energy

Before innovation, the binding constraint is how existing resources are used. On capital the diagnostic is blunt: unproductive firms are not competed out and productive firms do not grow. Where an American firm surviving to age 40 is about eight times its birth size — "up or out" — emerging-economy firms too often "flat and stay."

Talent is misallocated just as badly. As economies develop, the number of jobs demanding specialized skills rises many-fold, yet half the labor force is often locked out; when the United States was middle-income in the 1950s, 97 percent of doctors and lawyers were white men, whereas women and minorities now hold 52 percent of those professions, leaving the economy about 50 percent larger than the discriminatory pattern would have allowed. Energy is used no better: middle-income economies run energy intensity about two and a half times, and emissions intensity about three and a half times, that of high-income economies — gaps about using energy better, not switching fuels. Aggregated, if China used energy as efficiently as the United States, a Chinese worker would be about 80 percent as productive as an American, not 20 percent.

Switching models: investment, infusion, innovation

The escape route is a sequence, not a leap. Between investment and frontier innovation sits "infusion" — adapting global ideas to local circumstances — the missing ingredient in most middle-income countries: low-income economies concentrate on investment, lower-middle add infusion, and only upper-middle combine all three. Korea is the exemplar, sending engineers to learn electronics from Japan, extending collateral-free credit to firms with foreign licensing agreements, outcompeting the Japanese in Asian markets, and — once license prices rose — switching deliberately into innovation. Brazil chose the opposite path, subsidizing domestic R&D and patents behind a tariff wall; but its patent office could not tell new ideas from recycled ones, so firms collected subsidies while productivity relative to U.S. workers fell about 40 percent — leapfrogging straight to innovation rarely works.

Underneath both stories is Schumpeter's trio of forces — creation, preservation, and destruction. Most middle-income countries have no shortage of creation; they have a preservation problem, with incumbents (often state-owned enterprises) holding down the new. Schumpeter saw incumbents colluding to block entry, Aghion and Howitt cast entrants as the heroes who push the frontier, and Akcigit and Bill Kerr show that incumbents too can be innovators — so the task is to weaken preservation and rebalance the forces.

Evidence at a glance

- **Slowdowns are structural.** A growth slowdown is about three times more likely in a middle-income economy than in an advanced one, and lasts longer where economic freedom is lower.
- **The convergence gap is frozen.** Middle-income per-capita income relative to the United States has stayed near six percent since 1980; only China has broken out.
- **Firms are stuck.** Unproductive firms survive and productive firms fail to scale — "flat and stay" rather than "up or out" — and in Ukraine the productivity–growth link, once U.S.-like, later vanished entirely.

From X-rays to MRIs: the policy agenda

Discipline incumbents, reward merit, capitalize on crisis. Stop diagnosing growth with the X-rays of firm size, market concentration, and income inequality; use the MRIs of value added, socio-economic mobility, and emissions intensity.

- **Discipline incumbency.** Do not vilify incumbents; harness their talent, capital, and market power to advance development, not to preserve it.
- **Reward merit.** Connect young, small firms to markets, mentors, and finance so the best can scale.
- **Capitalize on crisis, and free enterprise.** Shocks weaken preservation, as when Korea let go of protected conglomerates in the Asian financial crisis; and economic freedom can shift fast, as Vietnam showed after 2017.

The Economics of Creative Destruction

A lecture by **Philippe Aghion**, Kurt Bjorklund Professor in Innovation and Growth, INSEAD, London School

of Economics and Collège de France.

The bottom line. Long-run growth is driven by creative destruction, and at its core lies a contradiction: innovation needs the prospect of rents, yet yesterday's innovators use those rents to block tomorrow's. Managing that contradiction — through competition, education, and flexicurity — is what lets economies avoid the middle-income trap and become both more innovative and more inclusive, including in the coming age of AI.

A paradigm built on a contradiction

The neoclassical growth theory of the 1950s, in Solow's elegant version, explained growth through capital accumulation but ran into decreasing returns; it needed technical progress for the long run yet never said where progress came from. The answer is creative destruction — Schumpeter's idea that new innovations make old technologies obsolete. Beginning in 1987 with Peter Howitt, a new Schumpeterian paradigm was built on three ideas: growth is cumulative, each innovator standing on earlier shoulders; innovations come from entrepreneurs chasing temporary monopoly rents; and each new innovation displaces the old.

The paradigm contains a built-in tension. Innovators need rents to justify the effort, but once they hold those rents they are tempted to block new entry so they are not themselves creatively destroyed; regulating a market economy is largely about managing this contradiction. In the basic Aghion–Howitt model, the equilibrium growth rate rises with the productivity and size of innovations and falls with the interest rate, so growth is positively correlated with the rate of creative destruction. That is why the United States, with more R&D and more breakthrough innovation, grows faster than a Europe specialized in incremental improvement.

Two extensions: competition and firm dynamics

The basic model implies that only outsiders innovate and that more competition is bad for growth — counterfactual results that forced the theory to evolve. Adding *step-by-step* innovation, where a laggard must climb each rung before overtaking a leader, introduces an escape-competition effect: leaders innovate to stay ahead while firms far behind are discouraged. Because advanced countries have more firms near the frontier, competition becomes increasingly growth-enhancing as a country develops — and empirically the relationship between competition and innovation is an inverted U.

A second extension models firms as collections of product lines that grow through creative destruction on new lines and shrink when rivals displace them. It predicts a highly skewed size distribution — many small firms, few large ones — and a positive relationship between firm age and size. That gradient is far steeper in the United States than in Mexico or India, because good U.S. innovators can reach venture capital and institutional investors and scale — so surviving long in the U.S. is itself a signal of quality.

Secular stagnation and the middle-income trap

The same tools illuminate two enigmas. In the United States, TFP growth surged between 1996 and 2006 and then declined. The surge coincided with superstar firms — Google, Microsoft, Walmart, Amazon — that harnessed the IT revolution; but weak competition policy let them expand unboundedly through mergers and new establishments, and entry has fallen since 2000. These leaders also stopped sharing knowledge, widening the gap to followers, discouraging laggard

innovation, and pushing markups up — so concentration rose and growth fell together.

The middle-income trap has the same structure across countries. Growth can come from catching up with the frontier — the advantage of backwardness — or from frontier innovation, and the two require different policies. Catch-up rewards education, trade openness, technology transfer, and better management, as in China after Deng Xiaoping; frontier innovation demands graduate schools, basic research, the freedom to fail, patient funding, venture capital, and competition. Economies that fail to switch fall into the trap: Korea's chaebols and Japan's keiretsu blocked the move to frontier institutions, and Europe, after catching up by the late 1980s, could not harness the IT revolution, lacking venture capital, patient funding, and a DARPA-style bridge between competition and industrial policy.

Evidence at a glance

- **Superstars, then stagnation.** U.S. TFP growth surged 1996–2006 alongside rising concentration, then declined as entry fell after 2000 and leader–follower technology gaps widened.
- **AI's potential is large but conditional.** Automating tasks in goods and ideas could add about 0.68 percentage point to annual growth over a decade — comparable to IT's 0.8 — provided competition is preserved.

Toward growth that is both innovative and inclusive

You need not choose between innovation and protection. Flexicurity, education, and competition make an economy more innovative *and* more inclusive at the same time.

- **Flexicurity.** Denmark's model — income support and active retraining after job loss — makes creative destruction work better and protects workers' health, vital as AI accelerates churn.
- **Education that levels the field.** The probability of inventing rises steeply with parental income; broad-based, high-quality schooling reduces "lost Einsteins" and raises both innovation and mobility.
- **Competition, made dynamic.** Judge mergers by whether they stifle entry and innovation, insist on data sharing, and pair competition with DARPA-style industrial policy and deep financial ecosystems.
- **Harness AI deliberately.** Preserve competition and open source to capture its growth potential, and pair strong basic education with active labor-market policies to capture its jobs potential.

The Economics of Dynamism: Firm Productivity, Entrepreneurship, and Industrial Policy

A lecture by **Ufuk Akcigit**, Co-Director, Growth Academy; Arnold C. Harberger Professor of Economics, University of Chicago (on leave); Deputy Chief Economist and Director of Private Markets, World Bank Group.

The bottom line. Every policy is a choice that moves resources from one use to another, so the question is never simply whether firms grow but whether growth is organic and productivity-driven. Middle-income economies stall because unproductive firms are kept alive, transformative entrepreneurs cannot scale, and blanket subsidies buy job numbers instead of efficiency. Better data lets policy target performance, not size, and reward the few firms that truly transform.

Growth accounting: the source of growth matters

We want economies to perform better, but spending money does not guarantee it. In a general-equilibrium world, subsidizing one firm means taxing another, so a policy that looks wonderful when we study only the firm receiving the dollar can leave the economy worse off — and most interventions lack the evidence to justify the trade-off. Decomposing growth into capital, labor, and efficiency, the efficiency term (the Solow residual) is the part we cannot see, which is why politicians prefer visible roads to invisible laboratories.

That preference is costly. Switzerland spends heavily on invisible R&D and labs, giving matching subsidies to universities so firms and universities advance together — an ecosystem, not a line item. Decompositions make the stakes concrete: in Korea and Poland, capital, labor, and efficiency all contributed positively; in South Africa about half of efficiency growth was negative, so whatever was gained was lost — the signature of a stalled economy. Plotting efficiency against income confirms it: Korea, Poland, and China grew as efficiency rose and stalled as it faded.

Static efficiency: are good firms allowed to scale?

Before creating new technology, an economy must use existing technology well. Aggregate productivity is the sum of each firm's productivity weighted by its market share, so the same technologies yield very different outcomes depending on who operates at what scale. The test is *responsiveness*: does a firm that receives a positive productivity shock actually gain resources and employment? In planned economies the correlation is absent; comparing West Germany with the United States, the responsiveness coefficient is about twice as high in the U.S., so the same innovation produces more efficiency there simply because input markets are less rigid.

Policy often makes this worse by keeping unproductive firms alive. When East Germany's firms were privatized after 1990, unemployment jumped and social unrest followed; after the head of the privatization agency was assassinated in April 1991, contracts shifted from asking a price to imposing labor commitments. Firms hired to satisfy those commitments, creating about 25 percent excess workers, and firms with binding commitments grew employment by about 30 percent. Thirty years on, the East–West productivity gap still exceeds 20 percent, whereas Poland, with no rich neighbor to rescue it, let restructuring happen and converged better.

Entrepreneurs, finance, and the design of policy

Not all entrepreneurs are alike. In Danish data, the likelihood of becoming an inventor or scientist rises sharply with ability, but the likelihood of starting a business falls with it — because most entrepreneurship is necessity-driven and subsistence. Distinguishing the roughly 10 percent of "transformative" entrepreneurs — identifiable early by whether they hire technical talent — from the 90 percent of subsistence founders matters enormously: transformative firms keep growing for a decade, while subsistence firms stay flat. Transformative entrepreneurs, like inventors, tend to be

educated and high-ability, yet parental income still predicts who reaches college even where tuition is free.

This heterogeneity should drive policy sequencing. With very little money, spend it all on equal access to education; with more, add a venture-style program that screens and funds promising entrepreneurs; only with ample resources subsidize incumbent R&D — yet intuition jumps straight to the last. Blanket instruments are captured: large firms claim most R&D tax credits because they have the resources to. France's "gazelle" program instead conditioned support on performance — growth above 15 percent for two consecutive years — selecting firms already trying to grow. Finance ties it together: among small firms, the transformative ones are precisely those that borrow to grow, so a credit shock hits them hardest and leaves lasting scars.

Evidence at a glance

- **Target performance, not size.** Among low-productivity firms, about 23–24 percent claim the R&D tax credit but only about 1 percent the gazelle program; among high-productivity, high-growth firms, about 96 percent claim the gazelle.
- **Finance is decisive.** Small transformative firms depend most on borrowing to grow, so a rise in the interest-rate spread produces a downturn from which the economy never fully recovers.

Designing for dynamism

Fix all four tires. Functional economies are alike — education, allocation, technology creation, and finance all work; dysfunctional ones fail for different reasons, so the task is to find the specific blockage.

- **Chase organic growth, not induced.** Subsidies that buy job numbers evaporate when they end; sustain jobs by letting resources flow to productive firms rather than propping up zombies.
- **Target transformative firms.** Use information — who a founder hires, how a firm performs — to reward the few that scale, conditioning support on performance, not size.
- **Sequence policy to the budget.** Education first, screened startup support next, incumbent R&D subsidies only when resources are ample.
- **Protect finance for growers in crises.** Small transformative firms rely on borrowing, so safeguarding their credit access during macro shocks prevents lasting damage.

Day 2 • Tuesday, July 28, 2026 • Chicago

Building Startup Ecosystems in Developing Countries

A conversation with **Samir Mayekar**, Managing Director of the Polsky Center for Entrepreneurship and Innovation, University of Chicago.

The bottom line. A real innovation ecosystem is not one institution but a set of interlocking parts—capital, an industrial base, entrepreneurs, incubators, universities and national labs, government, and connectors—and no single piece is sufficient. For most cities and countries the binding constraints are early-stage capital and a mobile, risk-tolerant workforce. The task is to diagnose your genuine strengths, focus on one or two sectors where you can win, and build bridges to where the money already is—rather than trying to recreate Silicon Valley at home.

An ecosystem, not a single institution

Mayekar came to the university as a founder who had lived the full arc. With co-founders he turned a powder that makes lithium batteries last longer and charge faster into a company, winning a business-plan competition that delivered about $1.5 million in angel investment over a single weekend. The first plant went to Japan—with a Japanese partner—because the United States was barely present in that China-, Korea-, and Japan-dominated market; later factories were built in the Midwest. Over roughly ten years the company raised about $200 million, took the U.S. Army as its first scaled customer, onshored as the supply chain turned America-centric, and was sold to a private-equity firm this year.

That experience maps onto a simple framework (from the organization UIDP): a true ecosystem needs capital, an industrial base, entrepreneurs, companies, incubators and accelerators, government as a partner, research organizations such as national labs, and connectors. His own company drew on nearly every piece—a $75 million Series B led by a local fund of former Argonne investors; Midwest automakers as partners; a university incubator with lab space that also supplied student hires; a cross-disciplinary class where he met his co-founders; a small-business (SBIR) research grant as the first money in; and a local clean-energy competition. Without any one of these, the company might not have survived. The lesson: diagnose where your ecosystem's parts are strong and weak, and attack the weaknesses.

The binding constraints: capital and workforce

Chicago and the Midwest have real strengths—deep manufacturing muscle, a university-driven research engine, and the country's most diversified big-city economy, in which no single sector exceeds 15% of metro GDP. But the deepest weaknesses are early-stage capital and workforce fluidity. Chicago is a private-equity town; founders with a risky idea and no cash flow must raise in Silicon Valley, and since venture capitalists will not travel, the fix is to go to them—plant an office in San Francisco and send entrepreneurs there. The workforce constraint is cultural: only places like Kendall Square and Silicon Valley have the density of people who will join a startup knowing most fail and simply move to the next one.

Agglomeration and collaboration are the levers. Within four or five hours of Chicago sits, through the Big Ten research alliance, the largest research-and-engineering base in America. Mayekar's

"Third Coast Foundry," a San Francisco hub built with eight universities, sits within five minutes of about $100 billion of venture capital; a recent demo day drew 250 investors managing $200 billion to 40 deep-tech startups. Government's highest-value role is a focused, patient bet, not picking companies: because the University of Chicago drew roughly half of federal quantum-research dollars, the governor placed a $500 million line item into a quantum park on a former steel-mill site, catalyzing a roughly $10 billion, largely private build-out and drawing DARPA and IBM—making Chicago the hub of the quantum economy. Such bets rest on years of research and must survive political cycles.

Evidence at a glance

- **Capital is concentrating in AI.** Per venture-industry (NVCA) data, about 90% of U.S. venture dollars over the trailing 12 months went to AI, with round and fund sizes ballooning; Mayekar expects a correction in roughly 18 months.
- **University equity models vary sharply.** U.S. universities typically take about 3–5% of a spinout (plus royalties), whereas a statutory 50% claim deters founders. The university's new $25 million fund writes $250,000 pre-seed checks into deep-tech spinouts, its Silicon Valley manager required to seat a partner in Chicago.
- **Seeding fund managers works.** Illinois committed state dollars as 10–20% of emerging managers' funds, helping create 30–40 new Chicago venture capitalists—the most diverse batch in the country, roughly half of them women.
- **Access is the real barrier.** About 60% of U.S. "unicorn" founders are first-generation immigrants; venture capitalists almost never fund cold outreach—warm intros through a small circle of schools dominate.

Policy implications

Focus, then connect. For most regions the fastest path is to pick one or two sectors with a genuine right to win, back professional early-stage investors, and bridge to global capital—not to clone Silicon Valley.

- **Seed professional investors, not companies.** Governments and universities should not pick winners; put capital behind independent managers and accept that most bets fail.
- **Build bridges to capital.** Send delegations and offices to where the money is, and cultivate a credible diaspora investor to anchor a country-focused bridge fund—as one did by helping U.S. firms enter Japan on condition they take its check.
- **Prize velocity over paperwork.** Public funding is throttled by compliance built for the rare fraudster and by cycles too slow for fast-moving technology; celebrate second chances, since failure culture changes only when people in power stand publicly with founders who have failed.
- **Concentrate public money where early dollars are hardest.** Science-based innovation—biotech, energy and materials, quantum—needs it; even early AI was seeded by public research.

Fertility Decline and Economic Policy

A lecture by **James Heckman**, Henry Schultz Distinguished Service Professor of Economics and Public Policy and Director of the Center for the Economics of Human Development, University of Chicago.

The bottom line. The statistics that drive fertility policy—above all the total fertility rate—are point-in-time snapshots projected across a lifetime, and they systematically mislead. Fertility has fallen almost everywhere over the past two decades, driven far more by the changing status and autonomy of women than by any single national policy. But the economic alarm is largely misplaced: in steady state a smaller population barely changes per-capita income, and the real costs are transient.

Measure fertility before you legislate

The crucial first question is how fertility is measured. The total fertility rate is a *period* measure: it sums age-specific birth rates observed in a single year and projects them as if one woman lived through all of them, assuming a stationary world across a reproductive span of more than thirty years—an assumption that never holds. The meaningful measure is the *cohort* rate—the number of children women actually bear over their lives—but it requires waiting, so demographers approximate it with duration and hazard models and "tempo versus quantum" adjustments. Replacement is roughly a total fertility rate of 2.1.

This distinction is not academic; it determines whether a policy looks like a success. Hungary's very generous program—no income tax for anyone with three or more children—appeared to work on the misleading period measures, yet examined carefully it did not raise completed fertility. France's storied pro-natal policies tell the same story: what looked effective has faded, marking "the end of the French exception."

A global, converging decline

Fertility has fallen across almost every region over the last 20 to 25 years, and the striking feature is convergence. The declines are most dramatic in East Asia, including China; South Korea now sits around 0.72—far below anything that could be called replacement—with the age at first birth near 32 and childlessness approaching 40% by age 35. India is below replacement, with urban fertility around 1.5 and rural around 2.1 as of 2024; rural India at replacement challenges the standard Western explanation built on the career opportunities of educated women. Iran fell rapidly after the Islamic Revolution; sub-Saharan Africa remains high; Europe is very low and alarmed.

The common thread is the changing status of women: more education, careers, and autonomy; postponement of the first birth; and rising childlessness (about 29% of Japanese women are childless at 40). Because the decline runs across every education level—not only among educated women—single-cause "opportunity cost" stories are too shallow. Culture transmits: Brazilian regions reached by soap operas portraying small, modern families saw fertility fall, a result later replicated in China. And across countries, fertility is higher where men take a larger share of childcare—near zero in Korea, where fertility is near zero. Notably, China's one-child policy is *not* the main cause: fertility kept falling after it was abolished and tracks a common East Asian trend.

Should we worry? Steady state versus transient

The economic case for alarm rests on confusing transient with steady-state effects. The transient burden is real—an older population strains pensions and public support during the demographic transition. But in steady state the penalty is small: adjusted-support-ratio calculations show per-capita consumption barely moving across total fertility rates from 1 to 3, with an "optimum" near 2.25. Heckman rejects the claim that productivity growth depends on the size of the population; productivity comes from ideas, research, education, and institutions, not headcount—and a recent

paper finds no link between population and economic growth. Indeed, an older, higher-saving population raises the capital stock, so a shrinking population can actually raise output per worker and welfare.

Evidence at a glance

- **South Korea is the outlier.** A national total fertility rate around 0.72, some regions near 0.32, childlessness close to 40% at age 35, and a first birth near age 32.
- **India below replacement.** As of 2024, urban fertility around 1.5 and rural around 2.1—rural India at replacement, contrary to Western expectations.
- **U.S. childbearing has fallen at every education level.** The share of women with four or more children fell from 38% to 18% (high school or less), 24% to 12% (some college), and 14% to 6% (college and above).
- **Steady-state economics is reassuring.** Per-capita consumption ranges only about 0.49 to 0.55 across total fertility rates from 1 to 3; the burden is transient, not permanent— much as, historically, the Black Death raised survivors' incomes when labor fell but land was fixed.

Policy implications

Don't legislate off a snapshot. Fertility is a long-run, non-stationary process; policy built on point-in-time total fertility rates will misdiagnose both the problem and the success of any remedy.

- **Use cohort, not period, measures.** Evaluate policy against completed cohort fertility and duration models; Hungary's apparent success dissolves under proper measurement.
- **Target the transition, not the steady state.** Concentrate resources on the transient burden—pensions and elder support during the demographic transition—since long-run per-capita income is largely unaffected.
- **Do not expect pro-natal subsidies to reverse the trend.** Generous cash and tax incentives have not durably raised fertility; the drivers are education, autonomy, and culture, and where policy does have leverage, fertility is higher where men share childcare.
- **Treat migration as a national, not global, offset.** It can rebalance a shrinking national population but cannot change the world total, and migrant fertility converges downward.

Actionable Evidence, Scalable Ideas: What Every Policymaker Should Ask

A lecture by **John List**, Director of the Becker Friedman Institute; Kenneth C. Griffin Distinguished Service Professor in Economics and the College, University of Chicago.

The bottom line. Before backing any policy, ask two questions: is the evidence **actionable**—does it establish cause, not correlation?—and is the idea **scalable**—will it survive outside the pilot? Most organizations get both wrong. Three "deception traps" come from broken comparisons you can design away; three more survive, and grow, when you scale. The frontier is a single field experiment that establishes cause, identifies who is affected and why, and tests whether it will generalize.

Two questions, and the literacy to answer them

Across firms as varied as United Airlines, Uber, Lyft, Meta, Walmart, and Anthropic—and governments from the United Kingdom to the White House to the Dominican Republic—every organization faces the same two questions daily: is the evidence actionable, and is the idea scalable? Almost all get it wrong. The root problem, List argues, is that schools do not teach *causal literacy*; teaching it would carry an enormous return, especially amid an AI-driven deluge of data and easy comparisons.

Six deception traps

The first three traps are broken comparisons that can be defeated by design. **Selection** appears when groups chose their own status: at Uber, the CEO noted that loyalty members spent about 400% more than non-members—but members were already higher spenders, and once that choice was accounted for the program was not incremental, in fact negative, because surplus flowed to inframarginal customers. **Hidden third variables** drive spurious links, as when ice-cream sales and drownings both rise with hot weather. **Before-and-after** comparisons where everyone gets the treatment mislead: Australia's plain-packaging rule coincided with smoking falling from 19.5% to 18.3%, but the full time series shows a pre-existing downward trend—the policy did not bend the curve.

The constructive principles follow: a natural experiment creates "statistical twins" when assignment is by luck or nature, and a field experiment creates them by controlling the assignment mechanism. Treating Uber's "bad trips" as good-as-random showed they cost about 6 to 9% of a rider's future revenue over 90 days; an apology recovered roughly 35% of that effect—but only when it carried a real cost, since words alone are cheap talk. The second three traps are *correct* comparisons that teach the wrong lesson, and they grow at scale. Under **Goodhart's law**, a third-grade math score predicts adult earnings, yet paying children to raise the score shifts the distribution without improving life outcomes—the mechanism (parents, teachers, schools, neighborhoods) is what matters, and AI benchmarks share the flaw. The **data deluge** guarantees spurious correlations amplified by confirmation bias, so "the data refiners are the new oil." **Survivorship bias** is the reflex to reinforce the bombers that returned where they were hit, ignoring those that never came back: you need the losers, not just the survivors.

From evidence to scale

Scaling adds two more failure modes: an irreplaceable *chef* will not scale where available *ingredients* can, and some inputs (great teachers) are hard to scale cheaply. The "voltage effect" is nearly a law—about 99.9% of ideas lose benefit-cost value at scale—so "move fast and break things" is art, not science. What transports across settings is *models*, not data: start from a simple model and keep generating data to find what is missing. At Walmart, 2022 price elasticities mislead today unless two situational features are pinned down—whether a price change crosses a competitor's price and an integer. At the Chicago Heights Early Childhood Center, the decisive

feature was teacher quality, so the pre-kindergarten program was built to work with mediocre, good, and excellent teachers, not only excellent ones. For governments, because a deployed policy is hard to claw back, staggered rollouts plus automatic sunset clauses let you learn what works while everyone eventually receives it—and stop what does not.

Evidence at a glance

- **Apologies, done right.** Uber bad trips cost about 6 to 9% of future revenue; a costly apology recovers roughly 35%, while cheap-talk apologies do not work.
- **Selection reverses conclusions.** Once selection is corrected, Uber's loyalty program was not incremental—even negative—because surplus went to inframarginal spenders.
- **Before-and-after misleads.** Australian smoking fell from 19.5% to 18.3%, but the pre-existing trend, not plain packaging, explains it.
- **Audiences differ.** In charitable giving, women respond more to social pressure and men more to price effects, so the ask should differ.
- **Nudges pay.** A Dominican Republic tax nudge raised about $185 million—roughly half a percent of GDP—in one tax season.

Policy implications

Design to learn. Treat every rollout as a chance to run a staggered, sunsettable experiment; the goal is not just a causal estimate but knowing who benefits, why, and whether it will scale.

- **Ask the two questions first.** Is the evidence actionable, and is the idea scalable? Answer both before committing resources.
- **Interrogate the comparison.** Diagnose selection, hidden variables, and before-and-after traps; prefer natural and field experiments that create statistical twins.
- **Distrust proxies, the deluge, and the survivors.** A good predictor becomes a bad target once you optimize it; spurious correlations multiply with data and AI; and inference needs the losers, not just the survivors.
- **Expect voltage drop, and design to learn.** Assume benefits shrink at scale, and identify the situational features—teacher quality, competitor prices—that determine whether a result transports; in government, use staggered, sunsettable rollouts so effects can be tested.

Day 3 • Wednesday, July 29, 2026 • Chicago

The Globalization of Entrepreneurship and Venture Capital

A lecture by **Emanuele Colonnelli**, Joseph L. Gidwitz Professor of Finance and Entrepreneurship, Booth School of Business, University of Chicago.

The bottom line. Venture capital has grown massively and gone global, yet its flows remain strikingly unequal—and in the poorest regions, overwhelmingly **foreign**. When high-growth founders are made to choose, they demand one thing above all: flexible **equity**, ideally from investors with local roots. Because local equity is scarce and access runs through personal networks, the policy task is to **build a resilient local equity ecosystem** and widen who can reach it—not merely to attract more capital of any color.

A global market that is still deeply unequal

Private capital—venture capital and private equity—has become one of the most closely watched engines of economic transformation, yet academics still have little concrete guidance for governments trying to grow a local ecosystem. The model is simple: passive limited partners (pension funds, foundations, universities, sovereign wealth funds, development finance institutions such as the IFC, and family offices) commit capital to funds run by general partners, who select and support portfolio companies until an exit through sale or IPO, with about 20% of returns accruing to the manager. The focus here is the risky early end—young, innovative startups with few assets—not established firms that rely on commercial banks.

That end of the market has exploded. Global commitments now dwarf the 2000 dot-com peak, even as a share of stock-market capitalization, driven by a handful of enormously valuable private companies, and more capital is allocated outside the United States. But growth has not brought equality: relative to GDP, poorer economies are even more underrepresented than their income alone would predict, and activity concentrates in a few urban hubs, with the United States and China dominant and London following. As the ecosystem globalizes it also changes shape—capital is becoming more local (China moved from a foreign-dominated to a locally dominated industry), and outside the U.S. and Europe the activity is almost entirely venture and growth equity rather than buyouts.

What entrepreneurs actually demand

To move past aggregate patterns, the lecture turned to a study of Africa—the most frontier region for this ecosystem, and a useful mirror for any country building one. Because data on unregulated private firms is poor, the research built its own, combining deal data and LinkedIn records with a cross-country IFC survey of about 4,500 respondents across 51 countries. The typical respondent is a young firm, roughly three to four years old with a median of six full-time employees, seeking about $500,000 and unable to raise a large share of it. Finance ranks near the top of every obstacle these founders name.

But simply asking founders is unreliable, because capital always arrives bundled with an investor. So the study ran an experiment: entrepreneurs ranked realistic but fictitious profiles that mixed the size, form, and source of an investment, with incentives to answer truthfully. The result was

unambiguous. Founders display a strong preference for **equity**, are averse to giving up control (board seats and large stakes), and place a modest premium on local knowledge. Strikingly, they barely respond to the identity of the investor—foreign versus local, government versus private—and non-financial support such as mentorship and networks does not move them. What they value in equity is its flexibility and the aligned incentive it gives an investor to help; rigid commercial-debt contracts do not fit firms with volatile revenues. This preference is constant across Africa and strongest among the largest firms.

Where the money actually comes from

The supply side tells a different story. About 80% of venture and growth-equity capital going to African companies is foreign—higher than in other emerging markets and far higher than in developed ones. China's footprint in Africa is mostly infrastructure lending, not venture; the flows come chiefly from the United States and Western Europe. Capital goes to highly educated, well-connected founders, and it matches on origin: U.S. money to founders tied to the U.S., French money to founders tied to France. Where founders demand local equity but supply is foreign and network-gated, the risk is an ecosystem with fewer local companies and wealth that accrues abroad.

Evidence at a glance

- **Concentration is extreme.** A few African hubs—Lagos, Cape Town, and Nairobi among them—account for about 40% of dollar activity while representing roughly 10% of GDP.
- **Founders are foreign-connected.** Roughly 80% of capital to African firms is foreign; about two-thirds of funded founders studied or worked outside Africa and about half are white—Latin America and India show high but lower shares, making Africa an outlier.
- **Connections predict access, not performance.** Colonial and language linkages predict who gets funded, yet foreign-connected firms do not outperform on IPOs, acquisitions, or employment growth—and after the 2022 macro reassessment, global investors pulled out hardest from the markets they knew least.

Implications for policy

Grow local equity, and widen access to it. The binding constraints are the scarcity of flexible local equity and the network barriers that determine who can reach capital at all—not a shortage of talent or a lack of mentorship programs.

- **Prioritize local, qualified investors.** Founders want equity from investors who can both fund and support them; cultivate a deeper pool of local equity providers, not just more capital.
- **Clear the regulatory path.** In many countries venture equity cannot sit on institutional balance sheets; reforms in Chile and Uruguay show how allocation can be eased.
- **Address exit risk, and counter the risk narrative.** Investors fear getting money out more than in; a fund-of-funds cushion (as with India's $10 billion program) and hard repayment data (as the IFC has shown) both help.

Growing Entrepreneurial and Innovation Clusters

A lecture by **Josh Lerner**, Jacob H. Schiff Professor of Entrepreneurship, Harvard Business School; Co-Director, HBS Private Capital Project.

The bottom line. Venture-backed firms punch far above their weight in innovation, which is why governments everywhere are pouring money into them. But most public programs disappoint—not through corruption but through good intentions poorly executed. The durable lesson is to **set the table** (fix the legal, exit, and labor environment) before handing out cash, to **let the market lead** through matching funds, and to **be patient**: clusters take decades. Middle-income countries should tailor to their strengths and turn "follower" status into an advantage.

Why venture capital matters—and why governments chase it

Venture activity is both far larger and far more global than it used to be. In 2001 some 80% of investment was in the United States; today the U.S. share is closer to 30–40%, with China and India accounting for much of the growth—and reaching Western levels of venture activity at far lower GDP per capita. Governments have responded with enormous public expenditures: over $5 billion in subsidies for the TSMC facility in Arizona, roughly $3 billion for the Micron semiconductor plant in Gujarat (about 18 months of budget for India's premier technical institutes, on Raghuram Rajan's calibration), and over $50 billion in Saudi Arabia's AI initiative alone.

Why the enthusiasm? Because the innovative footprint of venture is out of all proportion to its size. In the United States only about one in a thousand new businesses receives venture funding, yet half of the companies that have gone public in the last quarter-century were venture-backed—and that subset accounts for roughly 90% of the R&D spending of recent public companies. The pattern holds abroad: venture-backed firms represent about 40–50% of such R&D across emerging markets, and in developing countries outside China they account for roughly 31% of patenting (around 40% in knowledge-intensive industries). Meanwhile corporate research has been a mixed bag, with firms retreating from basic science; estimates suggest venture is three to four times more effective than corporate R&D per dollar.

Why so many programs fail—and what works instead

Everywhere with a successful industry shows the strong hand of government early on—the U.S., China, Israel, Singapore. But many governments have spent lavishly with little return, Europe being the clearest example. Two failure modes recur. First, officials plunge in without understanding the dynamics and chase the trend of the moment: Queensland, Australia poured mining royalties into biotechnology labs, only to watch the good ideas leave for Sydney or San Diego, while the firms actually winning export awards were drone-and-software businesses with a genuine local advantage. Second, and less common, programs get captured. Japan built the world's fourth-largest venture industry on subsidies in the 1990s, then watched it vanish once the subsidies stopped, because tax and employment policy still made entrepreneurship hard.

The constructive agenda rests on three principles. **Set the table**: get the institutional environment right first—exit markets (second-tier exchanges with real rule of law), flexible labor markets, working technology transfer, sensible taxes, and reduced penalties for failure (Italy's old bankruptcy register discouraged risk-taking; Singapore's "Phoenix Award" celebrated meritorious failure). This work is unglamorous and resisted by vested interests, but essential. **Let the market lead**: use matching-fund requirements so participants stake their own money, and welcome foreign investors where

local experience is thin. **Iterate and be patient**: clusters take decades—the modern U.S. industry began in 1946 but did not take off until the late 1970s.

Evidence at a glance

- **Young firms drive net job growth.** Across developed and developing economies, firms five years old or less account for the bulk of net new jobs.
- **Venture is high-powered money.** It is estimated three to four times more effective than corporate R&D per dollar in generating innovation.
- **A tiny slice, an outsized footprint.** About 0.1% of U.S. businesses get venture funding, yet venture-backed firms are roughly 90% of the R&D done by recent public companies; at the frontier, Sequoia estimates about 65% of its companies end with only nominal returns.
- **Europe's gap is structural.** U.S. venture activity is about four to one versus Europe; the U.S. has around nine companies worth $1 trillion or more founded in the last 50 years, Europe none worth more than 100 billion euros.
- **Skills are about opportunity-spotting.** Two of a few hundred Harvard College entrepreneurship students became billionaires in their early 20s—success tracked adaptability and vision more than formal credentials (cf. Jack Ma).

Implications for middle-income countries

Tailor the model; make "follower" an advantage. There is no one-size-fits-all playbook—copying Shenzhen or Silicon Valley wholesale fails. Fit the strategy to your comparative advantage, your market size, and your ability to adapt proven ideas to local conditions.

- **Target where you have an edge.** Not everyone needs a foundational AI model; the application layer, biotech, or clean tech may fit better—and let real, if early, local promise guide the choice.
- **Small nations must think regional and international.** Domestic markets alone are too thin; the UAE's success came from attracting firms with "a foot in different camps," not from funding only nationals for a national market.
- **Draw on the diaspora.** Singapore and Israel leaned on expatriates abroad for capital and expertise rather than relying solely on home-grown investors.
- **Exploit the follower's advantage.** Emerging-market entrepreneurs succeed by adapting what worked elsewhere to local conditions—Careem localizing ride-share for the MENA region is one example.
- **Learn from documented cases.** Two decades of experience—successes and failures alike, now codified in OECD volumes and country studies—give policymakers far more to learn from.

Spatial Creative Destruction and Regional Development

A lecture by **Leonard Wantchekon**, James Madison Professor of Political Economy and Professor of Politics and International Affairs, Princeton University.

The bottom line. Science drives growth through two channels—productivity and state capacity—but in many developing economies **state capture**, where firms and elites effectively own the state, structurally blocks creative destruction, and more competitive elections can make it worse. The way out is **decentralized safe havens**: start innovation in low-rent places and sectors elites ignore, anchor it in frontier agricultural research, and let success spread outward. Market access and research infrastructure—not top-down master plans—are the missing links.

Science, productivity, and the roadblock of state capture

Following Abdus Salam's line that "development is science with some noise," the lecture argued that science is the ultimate driver of progress—the reason life expectancy has climbed from roughly 35 two centuries ago to the low-to-mid 70s today. Science works through two mechanisms: **productivity** (a country's technological competence, in Nelson's phrase, is the basic factor controlling it) and, often overlooked, **state capacity** (technology lets states deliver services—a passport in Benin can now be issued in a day—and lets citizens act for themselves). And, as the innovation-and-competition literature shows, creative destruction is more likely in a democratic, competitive environment, with democracy most growth-enhancing near the frontier.

The obstacle is political. As work by Aghion and co-authors, and by Acemoglu and Robinson, has shown, political elites can block progress when dynamic innovation threatens their interests. Data from Italy make the mechanism concrete: larger firms are more likely to be politically connected, and connected firms are older and less innovative, so incumbents entrench themselves through connection to deter entry. But much of the developing world faces something worse—**state capture**, where firms and militaries own the economy and suppress competition through policy and repression alike (the Gupta brothers' capture of a South African ministry is one example). A study of local elections in Benin distinguishes *indirect* capture (a promise of policy) from *direct* capture (appointing the bureaucrat who will implement it). Direct capture is the more prominent form, and, counterintuitively, it rises when elections are *more* competitive: fearing promises won't survive a change of government, firms buy the official who stays regardless of who wins.

Decentralized safe havens

If central power is captured, where can creative destruction begin? The answer is **decentralized safe havens**: pick one place or sector where the rents elites seek are low—so they ignore it—and where people are eager for innovation, prove the model there, and let it spread through social learning and local externalities of proximity, culture, and complementarity. The parallel is China's special economic zones, but with science, not extractive industry, at the core—development as externalities moving between regions, not a top-down plan.

The idea grew out of a personal shock. Returning to his mother's fertile hometown, the speaker found a once-thriving agricultural market dead—because a single bridge to market had collapsed, after which young people left for moto-taxi work in the city. The data confirmed the pattern: across Sub-Saharan Africa, a 10% increase in soil suitability is associated with a 2% increase in poverty, because the best roads sit on the worst land. Colonial roads reached mines, not fertile soil, and since independence countries have merely repaved those same north–south routes, leaving weak market integration.

Anchoring safe havens in research

Fertile river regions are natural safe havens—some support four cropping cycles a year—but they need a research station to unlock that potential. Research infrastructure endures: a French oil-palm research center founded in 1922 still shapes its region's development. Frontier science is feasible on the continent (Ethiopia's space program for remote sensing, and teff, a drought-resistant grain now recognized globally as a superfood), and where it takes hold it becomes hard to reverse. Brazil's EMBRAPA (1973) is the model: research placed close to local ecology, a talent pipeline through a U.S. university, and, by recent estimates, a contribution of between 42% and 78% of total factor productivity growth.

Evidence at a glance

- **Fertile land, more poverty.** Across Sub-Saharan Africa, a 10% rise in soil suitability is associated with a 2% rise in poverty—because market access, not soil, drives outcomes.
- **Competition can worsen capture.** In Benin, direct state capture (appointing implementing bureaucrats) is the dominant form and increases as elections become more competitive.
- **Research infrastructure pays off.** EMBRAPA (1973) accounts for an estimated 42%–78% of Brazil's total factor productivity growth; a 1922 oil-palm center still marks its region a century on.

From despair to a plan

Be part of the solution. The three ingredients of a working safe haven are transportation access, innovation, and decentralized governance—and where the economic prize is large, even limited state capacity need not stop it.

- **Build the missing anchors, with extension.** The gap is rarely roads alone but the agronomic research station that raises productivity—and, as with EMBRAPA, the technical assistance and cooperatives that must accompany the science.
- **Connect farms to markets.** Link high-productivity rural areas to cities and to Nigeria's roughly 280-million consumer market, and add logistics to cut waste.
- **Put AI to work.** It can forecast prices and cut waste in agriculture, speed learning from research stations to farmers, and—via a Wikipedia-style platform—recover African history and languages otherwise being lost.
- **Institution-build directly.** The African School of Economics (opened 2014) is adding engineering, a research hub, and a private "agricultural complex" of experimental farms and cultural tourism—a roughly $44 million venture.

Day 4 • Thursday, July 30, 2026 • Chicago

AI and Economic Growth

A lecture by **Ben Jones**, Gordon and Llura Gund Family Professor of Entrepreneurship and Professor of Strategy, Kellogg School of Management, Northwestern University.

The bottom line. Any claim about what AI will do to growth and inequality is really a claim about three things: the **share of tasks** machines take over, **how productive** they are at those tasks, and how strong the **bottlenecks** are that tie tasks together. Because the economy is dominated by the things we *cannot* yet improve, bottlenecks make transformative, singularity-style acceleration unlikely—but they also mean that very capable AI can push value toward the workers who still do the hard, un-automated tasks.

A framework: three forces behind any AI claim

Economists model production as a set of *tasks*. To make anything—bread, a car, a cancer treatment, a unit of energy—you must complete many tasks, and each can be done by labour or handed to a machine. Total output is not a simple sum but a kind of average of what you produce at every task, and the way you average matters enormously. Written as a generalized (Hölder) mean, one parameter, ρ, decides whether output behaves like an arithmetic mean (big wins at any task drag the whole economy up), a harmonic or minimum function (the weakest links dominate), or something in between.

That gives three forces to track. First, **how many tasks** AI can take over. Second, **how good** it is at the tasks it takes—like arithmetic, where a phone chip beats a human by hundreds of billions to one, or like driving, where it may only match us? Third, **how strong the bottlenecks are**—whether the tasks we remain bad at quietly hold the whole process back. Automation thus has a horizontal dimension (the share of tasks moved to machines) and a vertical one (how much more each machine produces); public debate fixates on the first, but the second and third do most of the real work.

Automation, abundance, and labour's share

Moving tasks to machines raises output per worker almost mechanically: as people are squeezed into a smaller set of tasks, more hands per task make each one abundant while machines cover the rest. The effect on *inequality* is subtler and cuts both ways. More automation shifts income from labour to capital, because a firm now pays machines for tasks labour used to do—the fear that runs from Marx to today's headlines. Yet Marx's prediction failed: despite two centuries of automation, labour's share stayed near two-thirds in the United States and France for a very long time.

The reason is the vertical dimension. When a machine becomes *extremely* productive, it floods the market and the price of what it makes collapses, so it captures little of the economy—think of search or email, given away for almost nothing because computing is so cheap. Value then flows to the tasks that stay expensive because we are still bad at them, and those are largely done by labour. So "a lot of automation" and "rising labour share" can coexist. The dangerous case is *mediocre* automation—AI only slightly better than a worker but far cheaper, which takes the job without making anything meaningfully cheaper; the recent decline in the U.S. labour share fits automation accelerating while machine quality has not risen much.

Evidence at a glance

- **Computers "crush" us where they win.** A person needs about a minute to multiply two four-digit numbers; a phone chip does about 500 billion such operations in that minute—not a little better, but ridiculously so.
- **Bottlenecks are stronger than infinity.** If output is a harmonic average, making *half* of all tasks infinitely productive and free raises total output only from 1 to 2—not to infinity, or even to 5.
- **Efficiency does not equal share of GDP.** Computing power per dollar rose by a factor of about 10^{17} over roughly 50 years while overall productivity rose about threefold; information-technology equipment is still only about 4.5% of U.S. GDP.
- **Where the money actually goes.** Meals away from home are about 12% of GDP and housing about 30%—sectors famously resistant to productivity gains (a clogged sink can cost $1,000 for a plumber).
- **The farm shows the pattern.** A combine harvester once brought in about 6 million pounds of corn in 12 hours, yet corn near the farm costs a few cents a pound while corn chips in the store cost about $4.50—the un-automated steps set the price.

Implications for jobs, skills, and developing economies

Judge the forces, not the hype. Rather than predict a single number, ask of any sector: what share of tasks can AI do, how good is it at them, and where are the bottlenecks? That is what separates ordinary disruption from a genuine structural break.

- **Expect skill-biased pressure, with a twist.** Automation tends to replace lower-skill and routine work first. But because AI is a *cognitive* technology, physical bottleneck jobs (the plumber, the electrician) look relatively safe, while routine white-collar work is more exposed.
- **Baumol's cost disease still rules.** As agriculture and manufacturing were automated, their GDP shares fell and services rose; real incomes climb because everything else gets cheap, not because the service itself speeds up.
- **Speed brings its own risks.** Near-2% growth is ordinary disruption the U.S. economy handles well; a jump to 5% would be ahistorical and could provoke resistance, especially where entrenched incumbents block change—a danger for emerging markets.
- **Two-sided prospects for poor countries.** Automation is investment, so weak governance makes adoption hard and could widen divergence; but AI's latency means server farms need not be local, so—like mobile phones—cheap, even open-source AI could be accessed from afar as a broad enabler.
- **Innovation faces the same bottlenecks.** Telescopes and regression software far out-produce human eyes or hand calculation, yet conceptual and data bottlenecks make an R&D "singularity" doubtful—even as useful acceleration in health and technology remains possible.

A Master Class on Business Dynamism

A lecture by **Ufuk Akcigit**, Co-Director, Growth Academy and Arnold C. Harberger Professor of Economics, University of Chicago (on leave), Deputy Chief Economist, World Bank Group; with **Furkan Kilic**, Postdoctoral

The bottom line. Economies cannot be run as physical experiments, so the way to design credible policy is to build a quantitative "laboratory": high-quality **microdata**, disciplined by **theory**, calibrated to match a real economy, and then used to run counterfactual **policy experiments**. Because any single fact can support many competing stories, only a model that confronts *all* the symptoms at once can identify the true cause—and get the policy right.

The economist's laboratory

Chemists have a lab and physicists have their accelerators; economists have neither. We cannot spin up ten copies of Germany to try ten different policies. The substitute is economic theory: build a model whose firms and inventors behave like the real ones, and it becomes a laboratory in which counterfactual experiments can be run many times over, as one would replay a city-simulation game. The workflow runs from microdata to structural modelling to computational policy analysis, and every assumption bears directly on the conclusion—so the theory must be built on reliable data.

Microdata are essential because aggregates hide what matters. Two countries can post identical headline growth yet differ completely underneath: for the same innovation, the United States generates about twice as many jobs as Europe, and roughly half of all growth comes not from innovation itself but from *reallocation*—resources moving from less to more productive firms—a process far slower in some economies. Frictions live at the micro level, so policy must be designed there too.

Symptoms of declining dynamism

Read together, U.S. microdata since 1980 describe an economy losing dynamism. Market concentration has risen, average markups (price over marginal cost) have climbed, the profit share of GDP has roughly tripled, and labour's share has fallen. These are connected: within essentially every sector, higher concentration is associated with a lower labour share, so weaker competition leaves labour on the losing side. Meanwhile frontier firms keep advancing while laggards fail to keep up, the firm entry rate has trended down, the employment share of young firms (under five years old) has shrunk, job reallocation has slowed, and the dispersion of firm growth rates—a clean summary measure of risk-taking—has narrowed.

The danger is misdiagnosis. Any one symptom is consistent with many explanations, and a single data point can be fitted by many lines. Competing single-cause stories—lower corporate taxes, incumbent subsidies, changing entry costs, weaker knowledge diffusion, low interest rates keeping "zombie" firms alive, ideas getting harder to find, declining union power—each imply a different remedy, and treating the wrong cause is costly. Like a doctor, one must weigh all the vital signs before prescribing.

Evidence at a glance

- **We know how many people, not how many firms.** A quick search agrees the world holds about 8.2 billion people, but there is no agreed count of firms—a gap the World Bank is closing with a firm-level database spanning 64 countries over 1995–2025, including panels for 17 of them.

- **Micro firms are everywhere; the giants once were small.** Firms with fewer than five workers dominate the count, yet 30–40% of today's large firms (over 100 workers) were smaller a decade earlier—and were already more productive years before they grew.
- **Young firms create the jobs.** Most *net* job creation comes from young firms, and roughly 5–10% of firms ("high-growth" firms) create 50–60% of jobs—in all sectors, from textiles in Ethiopia to food and beverages in Hungary, not only in technology.
- **Jobs are often not durable.** In one Middle-Eastern economy, 55% of new matches dissolve within a year and 80% within five; of those who leave in year one, 34% exit the formal labour market—a "slippery jobs ladder" where churn does not lift wages as it does in advanced economies.
- **Estimates you cannot guess.** Calibrated to data, "expansion efficiency" is about 0.19 in the United States versus about 0.5 in India—a fourfold gap only structural estimation reveals.

From diagnosis to policy

Test every story in the model. Feeding each candidate cause into a calibrated model shows that most fit only one or two facts: declining corporate taxes explain the growth-rate dispersion alone, and R&D subsidies little more. The channel that matches all the moments—qualitatively and quantitatively—is *weaker knowledge diffusion*.

- **Diffusion, not just invention.** Creating an idea is one thing; spreading it through the economy is another. Killer acquisitions, patent barriers, and incumbents that deter entrants all reduce the diffusion of knowledge, which the model identifies as the core problem.
- **Put talent first.** With the same budget, subsidizing R&D, subsidizing schooling, or opening more university slots allocate talent very differently. R&D subsidies look effective at first (firms simply buy more equipment), but beyond about year eight the strongest policy is to strengthen the talent pool through education—especially by giving talented but poor children access to school.
- **Capability determines where subsidies work.** Estimated for Germany, the best R&D policy for Eastern citizens is to subsidize the more capable *Western* firms and let knowledge diffuse—a euro given to a low-capability firm is largely wasted.
- **Retire legacy policies.** The first federal U.S. R&D tax credit dates to 1981 (Minnesota followed in 1982); such temporary measures are renewed indefinitely for political reasons. Shutting down what no longer works frees resources for what does.

Skipping the Factory: Service-Led Growth and Structural Transformation in the Developing World

A lecture by **Fabrizio Zilibotti**, Tuntex Professor of International and Development Economics, Yale University.

The bottom line. Today's developing economies are leaving agriculture for **services**, largely skipping the factory stage that defined earlier transformations. This is not automatically a symptom of failure: careful measurement shows genuine **productivity growth** inside local consumer services. But the gains are **unequally shared**—they favour richer, urban households, while poorer, rural households gain most from productivity in tradable goods such as agriculture.

A different road out of agriculture

The canonical story of development is a sequence: labour leaves agriculture for industry, and only later moves into services. Italy ran this course—about half its workforce farmed in 1948, industry then expanded, and by 2019 services dominated—and the United States traced the same path a lifetime earlier. Plotted against income per capita, industrial employment follows a clear hump: it rises, then falls. Crucially, this decline is not mainly about China; Italy's industry shrank on the same trajectory the United States followed between 1880 and 1975, long before Chinese trade mattered.

The developing world today is taking a different route. As agriculture contracts, workers move overwhelmingly into services rather than manufacturing—"skipping the factory," or premature deindustrialization. Across countries poorer than China, the pattern is remarkably tight: for every 10-percentage-point fall in agriculture's employment share, manufacturing rises about 2.4 points while services rise about 6.4. India shows it, as do fast growers such as Ethiopia and Tanzania. The relationship holds for fast- and slow-growing countries alike, so it is not a mark of weak performance.

Do services deserve their bad reputation?

Services have long been treated as second-class. Goods are tangible and tradable; services are intangible and typically require producer and user to interact on site, which makes many of them local. A hierarchy runs from Marx's distinction between productive and unproductive labour, to the Soviet "material product" accounts that excluded health, education, finance, and personal services, to the modern view that manufacturing—with its scale economies, tradability, and clustering—is the true engine of productivity, so that industrialization *was* development.

The current transformation is dominated not by high-skill *producer* services (ICT, accountants, lawyers, designers) but by local *consumer* services—retail, personal services, health. Do they carry productivity growth, or signal its absence? Prices cannot settle the question, because observed service prices confound quality with efficiency and are often unavailable locally. The lecture's method sidesteps prices: since consumer services must be produced where they are consumed, household spending (demand) can be matched against local employment (supply), and the gap between demand-driven and actual employment growth reveals local service productivity. Applied through non-homothetic preferences to 380 districts in India and to Sub-Saharan Africa, it yields a striking result—productivity growth in local consumer services is real, and often higher than in tradables.

Evidence at a glance

- **The tight reallocation.** For every 10-point decline in agriculture's employment share across developing countries (1991–2019), manufacturing rose about 2.4 points and services about 6.4—most of the exit went straight to services.
- **Services out-grew tradables in productivity.** Estimated productivity growth in non-tradable

consumer services exceeded that in tradables in both urban and rural Sub-Saharan Africa—a pattern that also holds in India and, recently, China.

- **These are not merely reserve jobs.** Controlling for location and education, consumer-service workers have access to electricity and piped water similar to or slightly better than manufacturing workers.
- **Gains favour the better-off.** On a willingness-to-pay measure, rural poor households value tradable-goods productivity at about 16% of income (versus about 11% for the rich), while consumer-service gains are valued mostly by richer, urban households.
- **Place matters enormously.** Estimated consumer-service productivity growth reached about 11% a year in Bangalore versus about 2.4% in a poor rural district—so the rich in dynamic cities gain most, the rural poor from agriculture.

Policy implications

Treat services as productive assets, not consolation prizes. Service-led growth can be genuine growth. The task is not to force a return to the factory, but to raise the efficiency of services and connect them to the rest of the economy.

- **Understand where the productivity comes from.** Three channels appear to drive it: agglomeration externalities (denser consumer cities support more and more varied providers), the decline of informality (which lifts barriers to better jobs), and complementary technology and infrastructure—mobile phones, local banks, and property registries that support inventory, payments, and credit.
- **Do not abandon industrialization—but keep it grounded.** The benefits of industry remain real; the mistake is overambitious plans detached from a country's comparative advantage.
- **Build on revealed strengths and linkages.** Target clusters where an economy is already strong (for example, industries connected to tourism) and use linkage logic beyond the manufacturing input–output table.
- **Raise scale and efficiency in services.** Encourage the entry of more formal, larger providers rather than over-protecting small incumbents, as early Indian policy did.
- **Read the shift correctly.** What looks like deindustrialization is better understood as *slow* industrialization alongside a real service transformation—to be strengthened, not treated as failure.

Day 5 • Friday, July 31, 2026 • Chicago

The Paradox of Innovation Institutions: Designing Flexible Ecosystems that Accelerate Discovery, Invention, and Growth

A lecture by **James Evans**, Max Palevsky Professor and Director of the Knowledge Lab, University of Chicago.

The bottom line. Knowing and inventing are collective productions, and discovery rarely arrives by leaps. Researchers, firms, and whole countries move locally across a space of **adjacencies** set by the labs, skills, and partners they already have. That is why AI's default effect is to make search *more* conservative—but the same tools, deliberately repurposed, can map and rewire overlooked adjacencies, letting an innovation system reach discoveries it would otherwise never bump into.

Innovation moves through adjacencies, not leaps

Discovery is almost never the work of a lone mind. Scientists take others' data, borrow explanations, and hand predictions to agencies that can act. The natural unit of "complex power," then, is often the **national innovation system**—a set of research and development organizations whose relationships form meaningful adjacencies that are hard to jump-start where they are absent.

Modeling this space shows why. Placing a year's ideas, methods, and keywords into a high-dimensional manifold and letting it evolve by local (Brownian) motion predicts the next year's discoveries with about 95% discrimination. Progress is not random everywhere; it is *locally* random—a lab bumps into the method next door, not one across the space, because it lacks the equipment, supply chain, or partners to reach further. The rare exceptions matter enormously: Einstein's 1905 papers, built by reflecting on others' experiments, are so anomalous precisely because roughly 98% of new discoveries are driven by new methods and data. The improbable jumps are few, yet they are effectively the only ones that reshape a field.

Adding people to the model sharpens the point. Modeling researchers as random walks through the same space—who works with whom, on which problem, using which method—predicts *which* person will make a given materials-science discovery with a precision near 0.5, out of some 20,000 candidates, simply because they sit between a problem and its solution. Scaled to universities and countries, the same logic predicts whether a discovery lands in the United States, Germany, or China—so the task is not only to recruit great scientists but to **build systems that make the right adjacencies possible**.

The paradox: AI conservatizes by default

Large language models are less like individual minds than like **complex institutions**: cultural and social technologies that internalize the diversity and biases of the corpora they learn from. Reasoning models make this literal—trained by rewarding self-talk that reaches correct answers, they spin up on average three internal personas that debate, disagree, and prune bad ideas, and turning up a single "surprise" feature can double performance out of the box. A frontier system can be a society of hundreds of underlying models, which is why even its builders cannot fully anticipate where it will go.

The catch is that, left to their default use, these tools narrow rather than widen search. On average, people using AI fill in the missing cells of what is already known—a conservative move that shrinks the space of open questions as everyone converges on the same regions. Because models are tuned to be helpful, agreeable, and balanced, users surrender cognitive agency: on genuinely new problems they perform worse yet feel more confident. Yet the same technology, pointed differently, becomes a detector of the improbable. AI can identify where researchers are *not* looking—gaps created not by nature but by boundaries between fields and between the academy and industry—and simulate pathways that cross them, improving the discovery of energy-relevant materials and the repurposing of drugs, with effects larger still when validated experimentally with Argonne National Laboratory.

Evidence at a glance

- **Local motion predicts discovery.** A manifold of a year's ideas evolving by local randomness discriminates real from plausibly random next-year papers at about 95%.
- **Position beats brilliance.** Modeling people as walks through the knowledge space predicts the discoverer of a given material with precision near 0.5 among roughly 20,000 people.
- **Distance pays later.** A distant hit—like a new photovoltaic material—often underperforms nearby work at first but opens a neighborhood with a higher ceiling; a small British firm's polyester proved about as valuable as DuPont's entire space of nylon polymers.
- **Reasoning is a conversation.** Reasoning models generate internal dialogue at about 600% enrichment, averaging three debating personas rather than one.

Designing flexible ecosystems

Weaken memory; widen search. Trapped systems fail not for lack of talent but because incumbents, legacy institutions, and default AI all reinforce the paths already taken. The policy task is to keep idea and capital markets dynamic enough that new adjacencies can form.

- **Fund the young and the architectural.** Older scientists and mature firms tend to swap one component for another; young researchers and start-ups rewire whole systems across the economy. Unleash capital and idea markets for the disruptive bets incumbents avoid.
- **Do not import other countries' logjams.** Big consolidating teams and institutions run by past winners build up memory that weakens search—avoid replicating them.
- **Concentrate deliberately, then bridge.** Distributing narrow "centers of excellence" everywhere destroys the local adjacencies between disciplines and with industry; concentrate expertise regionally and connect it across sector boundaries, using AI-drawn maps of overlooked spaces to stage risky convenings that should sometimes fail.
- **Invest in your own problems.** Health and technology needs differ by country; imported solutions can misfire, as when fluid resuscitation for sepsis, based on a small 1940s study, was followed for 60 years before a global trial showed it was harming children and the under-40s.

AI for Development: The Case of Digital Agriculture

A lecture by **Michael Kremer**, University Professor in Economics and the College and the Harris School of Public Policy, University of Chicago; Director, Development Innovation Lab.

The bottom line. AI holds tremendous potential to advance development, but realizing it will rarely happen on its own: the biggest gains lie in applications where commercial incentives are weak, so **complementary public investment** is decisive. AI weather forecasts for smallholder farmers show the pattern—open, accurate, and cheap models deliver value only when paired with careful message design, trusted local data, and the institutions to reach farmers at scale.

AI for development: promise and the missing complements

Amid understandable concern about AI's risks, the more tractable question for this room is how to invest *now* to realize its potential across low- and middle-income countries—in small enterprises, education, health, and public-sector accountability. Frontier-lab models are important, but their benefits often require complementary investments where there is no strong commercial return. That is precisely where public policy has a role: either to incentivize firms to enter or to fund the work directly.

Digital agriculture illustrates the logic concretely, and AI weather forecasts illustrate it best. In rain-fed systems the central decision is *when to plant*. Planting just after the rains begin maximizes the growing season, but acting on a false start wastes seed and fertilizer and can destroy the crop. Farmers therefore hedge, plant late, and accept lower yields—a problem worsened by climate change, which erodes traditional rules of thumb, and by the historical neglect of tropical climates in meteorology. Evidence reviewed through the Development Innovation Lab's Innovation Commission is clear: across roughly half a dozen randomized trials and strong observational studies, farmers who receive forecasts change what and when they plant and how much they invest overall.

The forecasting leap—and the delivery problem

Cheaper, faster, open

An AI revolution has transformed the supply of forecasts. Many institutions now release global models as open access, so any country can download and adapt them; they run about 80,000 times faster than the supercomputer-bound dynamical models only a few centers in North America and Europe could once afford. In India, a blend of AI models—including systems from Google and the European Centre for Medium-Range Weather Forecasts—with a statistical model of historical rainfall predicted monsoon onset best, and adding a physics-based model did better still. Yet these systems are not yet optimized for smallholders, because the commercial return to serving them is small.

Making the message land

A good forecast is worthless if farmers cannot use it. Working with Precision Development and the Government of India, the team distributed *probabilistic* forecasts and tested how to phrase them—finding, for example, that "70% likelihood" was misread as 70 millimetres of rain until rewritten as "70 out of 100." Katie Cowell, a weather scientist with the Human-Centered Weather Forecast Initiative, stressed that this demands an interdisciplinary team, farmer focus groups, and trusted local observations. In Ethiopia, where mountains make forecasting hard, the group blends rain-gauge climatology the Ethiopian Weather Service trusts with an AI system it can now run on its own hardware, then converts grid-cell outputs into the sub-district boundaries farmers actually recognize.

Evidence at a glance

- **Farmers act on forecasts.** A study in Pakistan found farmers retimed applications around heavy or light rain; an Indian onset forecast changed timing, crop choice, and total investment.
- **Large, cheap returns.** A Benin study with intensive in-person training found gains up to $350 per farmer; when India scaled forecasts, the cost was pennies per farmer, implying benefit-cost ratios that could easily exceed 100.
- **Blends beat single models.** Combining AI and statistical (and then physics-based) models produced well-calibrated probabilities, with the blend outperforming each component.
- **Scaling worked.** Distributed widely for the first time in 2025, the forecast alone predicted an unusual monsoon pause; in 2026 the program reached 50 million farmers.

What it takes to scale

AI can do a lot—but only with complements. The forecast is the easy part. Trusted data, tested messages, delivery channels, and the institutions to sustain them are what turn an accurate model into higher productivity.

- **Use weather as the entry point.** Farmers check the weather constantly, so it is the natural hook onto which soil-chemistry advice, pest alerts, and other digital-agriculture services can be added over time.
- **Fund the public-good gap.** Multilateral development banks are moving in—the Asian Development Bank has announced $300 million for farmer weather forecasts and the Inter-American Development Bank a similar effort—and can spread lessons across countries and regions.
- **Guard against overfitting.** With many predictors and few once-a-year outcomes, models can look accurate yet fail in the field; governments procuring forecasts need transparency about fit rather than a simple "who predicts best" contest.
- **Reform procurement and evaluation.** Rules that focus on price suit hardware, not software, where quality matters; and rapidly changing technology demands new ways to evaluate impact.
- **Pair humans with tools.** As in structured-pedagogy teacher coaching, LLMs can help translate technical advisories into farmer-friendly messages, but adoption is highest when technology is combined with human contact.

Day 6 • Monday, August 3, 2026 • Washington, D.C.

Global Economic Update

A lecture by **Indermit Gill**, Chief Economist and Senior Vice President for Development Economics, World Bank Group; with **Ayhan Kose**, Deputy Chief Economist and Director of the Prospects Group, World Bank Group.

The bottom line. Potential growth has fallen for two and a half decades, and a run of shocks has left the world economy expanding below even that diminished potential. The post-pandemic rebound was strong but deeply uneven, convergence has stalled outside China and India, and a record wave of young workers is arriving just as investment and trade weaken. The realistic upside is **AI-driven productivity**—but capturing it requires trade openness, credible fiscal rules, and a far better business environment.

Growth is slowing, and the shocks keep coming

For a quarter century, growth rates have fallen steadily while the capacity to respond has narrowed under the weight of demographics, rising debt, and trade tensions. The decline in *potential* growth is rapid: for emerging market and developing economies it has slipped from about 6% a year in the 2000s to 5% in the 2010s and 4% in the 2020s. The world economy has absorbed a great deal and is still standing, but keep growth falling and eventually a single shock could tip it over.

Ayhan Kose framed the era as "one shock after another": the pandemic, supply disruptions, a sharp rise in interest rates, the Russia–Ukraine war, an inflation surge, and now conflict in the Gulf. That conflict is the largest oil-and-gas disruption by some measures in five decades, and its effects cascade—first through fuel, then through food and fertilizer as importers face higher prices and exporters impose bans, and finally through finance as higher inflation lifts policy rates. For food importers the stakes are stark: in Nigeria, the poor spend around 70% of household budgets on food. Because most developing countries are small economies, there is no substitute for more trade—even when larger partners turn inward.

An uneven recovery and a jobs wave

Convergence has stalled

The 2020 downturn was the deepest global recession since 1960 yet drew the strongest recovery, the world economy expanding about 14% over five years. But the rebound was driven by advanced economies. One-third of emerging market and developing economies—and close to 60% of fragile and conflict-affected states—still have lower per-capita income than in 2019. Stripping out China and India, convergence toward advanced-economy incomes has essentially stopped; in Sub-Saharan Africa, per-capita income has been roughly flat for 15 years. When per-capita incomes stop rising, Gill warned, the entire pathology of development changes.

The 1.2 billion

The defining challenge is demographic. Some 1.2 billion young people will reach working age between 2026 and 2035—the largest youth cohort ever—and the next surge will crest in Sub-Saharan Africa, South Asia, and the Middle East and North Africa, where East Asia's old path of manufacturing and trade is less available. Yet investment and trade growth both run near half their 2000s pace, and investment per worker has stagnated outside China. Growth that does not create jobs breeds political and social strain—so the task is not just more growth, but growth that employs.

Evidence at a glance

- **Potential growth is falling everywhere.** Advanced-economy potential has roughly halved since the 2000s; emerging-market potential has slipped from 6% to 5% to 4%.
- **2026 is subdued.** World growth is expected around 2.5%; emerging and developing economies near 3.6%—about 10% below potential, and only 2.6% excluding China and India.
- **Debt has surged.** Government debt in emerging and developing economies has more than doubled to about 77% of GDP over 10 to 15 years, leaving little room if interest rates rise.
- **Trade has not collapsed—uncertainty has spiked.** World trade rose from about 20% of GDP in 1945 to 60% in 2022; the ratio has since plateaued but not fallen, while trade-policy uncertainty has climbed.
- **Barriers start at home.** Middle-income economies impose higher tariffs and non-tariff barriers on each other than they face in advanced markets—and now account for more than half of world trade.

The policy response

AI is the credible upside—if the fundamentals are in place. In the World Bank's more optimistic scenarios, AI could deliver the fastest productivity growth in decades and reverse the slowdown, but only economies with open trade, sound public finances, and a strong business environment will capture it.

- **Liberalize trade unilaterally.** For a small economy, opening up yields large gains regardless of what the United States or other partners do.
- **Adopt and enforce fiscal rules.** Emerging economies have improved monetary policy but still lag on fiscal management; credible rules raise the odds of high-quality adjustment, as Turkey's earlier turnaround showed.
- **Fix the business environment.** With foreign direct investment near two-decade lows relative to GDP, improving conditions for firms is essential to attracting capital.
- **Invest in infrastructure and people.** World Bank Group's response centers on aggressive investment in physical, digital, and human capital, plus mechanisms to mobilize private capital.
- **Take AI seriously—especially if poor.** AI can lift agricultural productivity and the quality of health, judicial, and legal services; poorer economies are often better prepared than assumed, as small and informal firms already adopt digital tools—and faster growth, anchored in large economies whose spillovers lift others, is what gives the distribution debate anything to work with.

Industrial Policy for Development: A Global Overview

A lecture by **Tristan Reed**, Economist, Development Research Group, World Bank Group.

The bottom line. The debate is no longer whether governments should do industrial policy but how. Defining it as government action to grow a strategic business activity, the World Bank's new global report urges countries to match tools to three characteristics – government bandwidth, local market size, and fiscal space – and to reach first for low-cost "public inputs" that address market failures directly before turning to subsidies, tariffs, and other incentives.

Reframing an old debate

The starting point is the 1993 *East Asian Miracle* report, the World Bank's first policy research report. Written amid a fierce debate over the role of the state versus the market, it credited growth to macroeconomic fundamentals – human capital accumulation, high savings, macro stability, and modest fiscal deficits – rather than to industrial policy, and concluded that even where interventions had succeeded in Japan and Korea they were unlikely to be replicated because, in the words of the foreword, those countries had a unique culture, politics, and history.

This new report reopens the question from a global, rather than East Asian, vantage point, organized around five questions: what industrial policy is, who does it, how to do it, which activities to target, and how to get the institutions right. It defines industrial policy as government action to grow a strategic business activity – where "industry" means any business activity (agribusiness, critical minerals, professional services, tourism), "activity" captures specific tasks and segments of a value chain, and "strategic" signals that government has judged one activity more important than others.

Crucially, developing countries are not bystanders. A survey of the Bank's lead country economists found that 80% had been asked about industrial policy by clients in the past year, overwhelmingly motivated by economic development rather than climate or security. Reading national development plans shows poorer countries list *more* priority industries – reflecting a hunger for diversification – while tariffs run more than twice as high in low-income countries and subsidy spending is heaviest in upper-middle-income ones.

A toolkit matched to country characteristics

Rather than ask whether to intervene, the report starts with how. It sorts tools into three families: **public inputs** (industrial parks, skills development, market-access assistance, and quality infrastructure such as standards and metrology), **market incentives** (subsidies, tariffs, local-content requirements), and **macroeconomic interventions** (a competitive exchange rate, a general R&D tax credit). Public inputs and direct subsidies are "first choice" because they address market failures directly; tariffs and other indirect tools are "second choice."

Which of these is feasible depends on three country characteristics. *Government bandwidth* – a workforce able to engage firms and gather information – is needed to tailor public inputs. *Local market size*, measured by the middle class and any customs union, determines whether economies of scale can be reached behind protection. *Fiscal space* determines whether cash subsidies are affordable at all. A country short on all three still has options: industrial parks (which can be self-financing), commodity export bans, and exchange-rate devaluation. As bandwidth grows, the cheap first-choice public inputs come within reach; market size unlocks second-choice tools; only fiscal space opens the full menu of subsidies.

Evidence at a glance

- **Industrial parks can spur local development.** A *Journal of Development Economics* study mapping Africa's parks found, within a 10-kilometer radius, a shift out of agricultural jobs and gains in education and durable goods such as refrigerators, televisions, and mobile phones.
- **Targeted public inputs work when tied to demand.** Costa Rica attracted Intel by co-designing a one-year associate's degree with the firm's HR department and a local university; market-access assistance in Latin America and quality certification (the Benin Republic's pineapples) have lifted exports.
- **Well-designed subsidies can be both effective and efficient.** Korea's heavy-and-chemical-industry credit program raised sales, employment, and – after a decade of learning – total factor productivity, leaving the economy about 3% larger each year (Choi and Levchenko); Romania's software payroll-tax exemption helped make it a leading software-services exporter.
- **Innovation grants pay off through infusion.** Brazil's Embrapa turned the country from a net food importer into an exporter; its Finep grants let firms import and adapt foreign technology, with added payroll-tax revenue covering the cost.
- **Tariffs rarely succeed alone.** The 1890 McKinley tariff (70%) built a U.S. tinplate industry only over a decade, behind a huge domestic market and with domestic inputs – and Irwin argues it would have emerged anyway.
- **Export bans need market power.** Indonesia's nickel ban drew investment in smelting because it held a large global share; an identical bauxite ban failed because it did not.

What to target – and for how long

Target what is new, keep competition alive, and be patient. Because true market failures are hard to measure, the practical rule is to promote activities the economy has not done before – and to expect results over five to ten years, not one political cycle.

- **Favor new activities.** Learning comes from doing something new, yet Kenya subsidizes long-established tea and 92% of Serbia's subsidies go to firms that do no R&D.
- **Preserve competition.** Support dispersed across many firms raises productivity more than picking a single winner (Harrison and co-authors, China); where the domestic market is small, exposure to trade must substitute for domestic rivalry, as Japan's development-era export-disciplined cartels show.
- **Handle jobs and duration carefully.** Broad, easy-to-access adjustment programs (Brazil's reskilling) beat narrow ones (U.S. trade adjustment assistance); the hardest political problem is not letting losers go but letting winners go, as reversals of green industrial policy in the United States and Brazil illustrate.

Industrial Policy in the Digital Age: Lessons from Asia

A lecture by **Siddharth Sharma**, Lead Economist, Office of the Chief Economist for South Asia, World Bank Group; with **Alessandro Barattieri**, Senior Economist, Office of the Chief Economist for East Asia and Pacific, World Bank Group.

The bottom line. Asia is fast-growing but faces a productivity and jobs challenge that industrial policy has so far addressed unevenly. Drawing on the World Bank's April 2026 regional updates for South Asia and East Asia and Pacific, both speakers argue that targeted measures work only atop strong foundations – infrastructure, human capital, and open, undistorted markets – and that artificial intelligence is now reshaping which tools make sense.

South Asia: fast growth, a jobs gap, and an AI test

South Asia grew 7% in FY2025, again the fastest-growing emerging-market region, powered by domestic demand in India and a better-than-expected recovery in Sri Lanka. But it is exposed to conflict in the Middle East: a large net energy importer (energy is about 5% of the consumption basket, above the EMDE average), it also sends about 13% of its exports to the Gulf and hosts roughly 9 million migrants there whose remittances reach about 8% of GDP in Nepal. With crude expected to average $94 a barrel over FY2026, the risk has materialized on the upside.

Against this, 2026 is a year of trade opening. India's new free-trade agreements with the EU and UK cut tariffs on about 95% of traded goods and roughly double the share of its output covered by such agreements to 30%, near China's level; Sri Lanka's phase-out of para-tariffs would cut its average import duty by nine percentage points. First-order estimates show real incomes rising across the distribution – more for rural households, and, in Sri Lanka, more for the poor because manufactured food faces the largest cuts. Yet the jobs challenge is stark: about 16 million young people will enter the labor force each year for 15 years, against roughly 12 million jobs created annually over the past decade.

Artificial intelligence could help or hurt. About 22% of South Asian jobs are highly AI-exposed – below the EMDE average because so many work in agriculture – but those jobs earn about 42% of labor income, and more are complementary to AI (15%) than at risk of automation (7%). Job-posting data confirm it: hiring in highly exposed, low-complementarity occupations fell about 22%, while the export-oriented ICT sector is upgrading toward more AI-complementary skills.

East Asia: a productivity problem and three pillars

East Asia's recent growth outside China has come from capital accumulation, not productivity, and firm-level data show a widening gap between leading firms and the global frontier, especially in digital sectors. That makes the capital-driven, labor-intensive model risky – and makes AI, whose productivity potential is large but uncertain, worth taking seriously.

The East Asia and Pacific update frames industrial policy broadly, as three pillars. **Foundational public goods** come first: many countries sit below the global median on health and education, and human capital, not algorithms, is the binding constraint. **Addressing policy failures** – "do no harm before you do good" – means, distinctively for the region, opening highly protected service sectors; Vietnam's post-2007 services liberalization raised labor productivity about 2.9% a year over nine years. Only then come **targeted interventions** against market failures, from Indonesia's resource downstreaming to Vietnam's plan to train 50,000 semiconductor engineers.

A central distinction runs through the evidence: effectiveness versus efficiency. A policy can hit its target yet cost more than it is worth. Korea's 1970s heavy-industry program was both effective and efficient, with credit flowing to the highest-welfare sectors; China's shipbuilding drive lifted its global market share from 10% to nearly 50% but, by the study's cost-benefit measure, was inefficient. And along the AI value chain, developing East Asian economies sit mostly at the low-value assembly

end – a position Korea's own three-act semiconductor history suggests can change only by building capital and technology at home.

Evidence at a glance

- **Trade openness is rising fast.** India's FTAs cover about 95% of traded goods; Sri Lanka's reform would cut its 19% average import duty by nine points. The average real-income gain is about 0.2% in India but about 2% in Sri Lanka, which liberalizes across all partners.
- **South Asia leans on protection.** Industrial-policy measures nearly doubled after COVID; the most common type is trade-related import protection, versus subsidies in the average EMDE, and countries tend to target larger, higher-productivity firms.
- **Its track record is mixed.** Event studies show import-protection measures do cut imports, but export-promotion measures show no significant effect – consistent with limited fiscal space, thin government bandwidth, and missing complementary inputs.
- **Targeted subsidies need strong foundations.** Matching global firm and subsidy data, export promotion and direct transfers correlate with higher productivity only in G20 economies, not in developing ones.

Implications for policymakers

Build the foundations first; target only where market failures are clear. Broad-based reforms – infrastructure, business climate, institutions, open services – do most of the work; targeted tools should complement them, and should always be judged on efficiency, not just effectiveness.

- **Tilt growth toward jobs with first-best tools.** Where failures are well identified, favor industrial parks, skills programs, and market-access assistance – for example, to help the ICT sector adjust to AI-driven reshoring pressure.
- **Weigh the data-center rush carefully.** Because AI compute is internationally traded, location reflects fundamentals – reliable power, water, and, strikingly, governance – and spillovers are uncertain.
- **Do no harm before doing good.** Liberalizing protected services is itself a powerful industrial policy, often safer than new subsidies.

Day 7 • Tuesday, August 4, 2026 • Washington, D.C.

Industrial Policy in the Middle East and Africa: Challenges of Conflict and Structural Transformation

A lecture by **Caglar Ozden**, Deputy Chief Economist for the Middle East and North Africa, World Bank Group; with **Andrew Dabalen**, Chief Economist for the Africa Region, World Bank Group.

The bottom line. Both regions use industrial policy intensively, yet neither has delivered the structural transformation it needs. In the Middle East and North Africa, the tension is between an ambitious, sovereign-wealth-funded AI push and the drag of dominant state-owned enterprises and missing data. In Africa, fiscal constraints, thin complementary ecosystems, and short-lived support blunt the tools – so both speakers urge patience, capability-building, and honest diagnostics over "champions."

The Middle East: the good, the bad, and the ugly

The MENA region spans three worlds – high-income Gulf oil exporters, middle-income oil importers such as Egypt and Pakistan, and fragile, conflict-affected states – but the story rhymes across them. Governments treat industrial policy as a cure-all, yet many challenges are not market failures at all; they are more fundamental problems of governance and institutions. Framed as "the good, the bad, and the ugly," the diagnosis pairs a genuine AI opportunity against two structural weaknesses.

The bad is the state-owned enterprise. SOEs generate roughly 20% of the region's GDP and appear even in contestable sectors – wholesale trade, construction, chemicals, basic metals – far more than elsewhere in the world. From Saudi Aramco and Morocco's phosphate giant to Egypt's military-owned resorts, they enjoy preferential access to cheap capital, crowd out private firms, and dull contestability. As one discussant noted, they are systematically less dynamic: where SOEs expand, private R&D shrinks and entry falls, and because jobs there feel safe, they also draw high-skill talent away from entrepreneurship.

The good is artificial intelligence, where Saudi Arabia and the UAE are pulling far ahead – topping regional AI-preparedness rankings, building data centers and homegrown large models, and, in the UAE, planning for AI to run half of government procedures. Much of this is financed not by private markets but by sovereign wealth funds: Saudi Arabia's PIF invests the bulk of its portfolio domestically, in an explicit act of industrial policy. The ugly is data: even wealthy Gulf states score below comparators on statistical performance, information is often uncollected or unshared, and the Bank's own enterprise surveys capture almost no SOEs – leaving much policy advice closer to speculation than analysis.

Africa: many tools, little transformation

Africa has been an active user of industrial policy – special economic zones, export and import bans, tariffs, local-content rules – yet none of it has produced the structural transformation the region needs, even as the working-age population of Sub-Saharan Africa is set to nearly double over the next generation. The stakes are high and the headwinds severe: debt averages about

60% of GDP, with a fifth to a quarter of revenue going to debt service.

The tools African governments reach for are overwhelmingly border measures – reflecting constraints, not preferences. Tariffs raise revenue when fiscal space is scarce, but they protect imperfectly, penalizing firms that rely on imported inputs. Two further problems recur: support is too short and thin, when learning in manufacturing takes five to ten years of sustained backing, and complementary ecosystems are missing – a zone with reliable power is undercut when its suppliers outside the fence have none. The gaps are stark: all 47 small Sub-Saharan economies together consume less electricity than South Korea, and credit to firms is about 30% of GDP against over 100% in East Asia.

The way forward is a two-part diagnostic: is a policy *feasible* (fiscal space, bandwidth, market size) and will it be *effective* (given the ecosystem, the industrial base, and the realism of anchoring growth in resources)? For many countries the honest answer is to fix foundations first before attempting production subsidies, and to lean on regional integration to overcome small market size.

Evidence at a glance

- **Sovereign wealth is the region's industrial policy.** Unlike Norway's globally diversified fund, Saudi Arabia's PIF keeps most of its portfolio at home, channeling roughly 16–17% into information technology and communications and backing a homegrown AI model.
- **Same tool, opposite results.** Indonesia's export ban built a nickel industry but failed for bauxite – a warning for African countries tempted to ban lithium, cobalt, or copper without global market share or competitive costs.
- **Discipline separates success from rent.** Nigeria's open-ended cement tax holidays became rent extraction, while Rwanda let non-performers fail; Kenya's horticulture succeeded where Tanzania's did not because it supplied logistics and cold storage.
- **Capabilities are decisive.** Egypt's Red Sea tourism worked through land reform, new airports, and, crucially, training institutions for hotel staff – and Korea's rise drew on scientists recalled from abroad and universities funded 60–70% by industry.

Implications for policymakers

Diagnose honestly, build capabilities, and use industrial policy sparingly. Where the private sector is thin and ecosystems weak, foundations and learning matter more than any single instrument – and success should never rest on a named "champion," which can fail.

- **Attack the real constraint.** As one discussant pressed, the common denominator is fixing inefficiencies – externalities, spillovers, financial frictions – with performance-based, not size-based, support; capability-building is part of a broader ecosystem, not a substitute for it.
- **Invest in complementary ecosystems.** Reliable energy, logistics, skills, and credit will otherwise undermine any intervention, especially in Africa.
- **Open the data, and integrate the region.** In MENA, transparency is a precondition for good policy and better governance; in Africa, regional integration enlarges markets and competition, widening the feasible instrument mix.

Revisiting Industrial Policy in ECA and LAC: Strategic Options for Today

A lecture by **Ivailo Izvorski**, Chief Economist for Europe and Central Asia, World Bank Group; with **William F. Maloney**, Chief Economist for Latin America and the Caribbean, World Bank Group.

The bottom line. Industrial policy is back in fashion across Europe and Central Asia (ECA) and Latin America and the Caribbean (LAC), but in both regions it is rarely well planned, well funded, or evaluated. The evidence points the same way: growth is not held back by being in the "wrong" products, but by weak capabilities to adopt technology and place informed bets. The main engine of development is productivity — so market interventions should be used sparingly, disciplined by a real market-failure test, and built on the fundamentals of competition, capabilities, and education.

Two regions, one temptation

ECA converged rapidly toward EU income levels over the last 30 years — 11 formerly centrally planned economies have graduated, and the EU "convergence machine" pulled trade and investment forward. Yet as convergence advanced, population growth fell, aging accelerated, and the contribution of productivity dropped sharply. The paradox Izvorski underlined is that these economies lived through the planned era — de facto industrial policy *par excellence* — and still ended up productivity-constrained.

LAC arrives at the same debate from the opposite direction. It has a long, generally unsuccessful history of industrial policy, and it does not grow: Chile, which embraced the market-friendly model most wholeheartedly, is projected to grow about a percentage-and-a-half this year with little productivity growth in a decade. In both regions the renewed enthusiasm is driven from outside — announcements have surged since COVID, propelled by national-security concerns in ECA and by the sight of China and advanced economies intervening heavily in LAC.

Crucially, both insisted on the same discipline: a policy is justified only where there is a well-identified, structural market failure that reforms cannot fix. Absent it, intervention mostly protects incumbents — "the firms of yesterday" — against limited bandwidth, thin fiscal space, and a real danger of capture.

It is how you make it, not what you make

Where the money goes tells the story. In ECA, more than half of interventions target agriculture and food production, with very little on high tech and almost nothing on "moonshots" — no chips, no new electric cars — and green-transition bets such as Turkey's solar panels and Poland's heat pumps sit on established, not frontier, technology. Having gone from the world's lowest-subsidy region to one of its highest, ECA now spends more without spending better — and where programs were evaluated, from Kazakhstan's special economic zones to North Macedonia's innovation aid, the productivity gains were hard to find.

The capabilities constraint

Maloney reframed the whole question: LAC does worse than other regions in *most* of its sectors, so the puzzle is not the product but the capacity to place and execute informed bets. Development is fundamentally about people making risk-return calculations, and there is enormous heterogeneity

around any good — Norway's experience with gas is not Nigeria's. History drives the point home: many countries near 40% of U.S. income in 1860 — Spain, Japan, Korea, Portugal, Sweden — took off, while LAC did not, undone by a weak ability to absorb the second industrial revolution's technologies. Chile, the world's largest copper producer in 1860, nearly lost the industry by 1900 and recovered only when U.S. firms brought new technology, while Japan built Mitsubishi, Sumitomo, and Hitachi out of copper — the lesson being not to abandon a product but to master its technology.

Evidence at a glance

- **Spending rose, sophistication did not.** As subsidies climbed, productivity growth slowed and export sophistication in several ECA countries deteriorated over the same period.
- **Concentration and weak evaluation.** Turkey, Russia, and Poland account for about 70% of ECA GDP and dominate its industrial policy; where programs were assessed, Kazakhstan's special economic zones could not be judged and North Macedonia's innovation aid lifted jobs and wages but left productivity unchanged.
- **The frontier-distance trap.** Only firms near the technological frontier can "innovate their way out" of competition — about 50% in the UK and France, but only about 7% in Chile in response to the China shock.
- **Capabilities gaps are old and deep.** LAC lags on literacy and STEM graduates, holds about 0.5% of the world's top-1,000 universities, sits mid-pack on management quality, and has shallow financial markets.

Policy implications

Fundamentals first, interventions sparingly. Industrial policy is not the main engine of development — productivity is. Where it is used, it must clear a genuine market-failure test, stay within limited state bandwidth, and rest on competition, capabilities, and education rather than substitute for them.

- **Demand a market-failure diagnosis.** Keep the failure analysis as a disciplinary device; identify what is structurally broken before spending, and prefer tailored public inputs to market interventions.
- **Build capabilities and fix the plumbing.** Engineers, designers, and managers — not just financial incentives — decide whether a bet succeeds (Brazil's Petrobras tanker initiative failed for lack of them), while shallow credit and slow contract and bankruptcy resolution (about two years in much of LAC versus three months in the U.S.) keep entrepreneurs from placing big, diversified bets.
- **Learn the right lesson from Korea.** Its success rested on massive human-capital investment and structural reform, not industrial policy alone — so preserving market mechanisms and competition must come first.

Learning from Lending

A lecture by **Peter Henry**, Senior Fellow, Hoover Institution and Dean Emeritus, New York University's Stern School of Business.

The bottom line. Sixty years of development lending hold lessons we have barely used. A "dual hurdle" framework asks which public investments clear both a social threshold and a private one — so that private capital can be mobilized where projects pay for themselves, freeing scarce public money for projects that only serve society. Applied to Jamaica's post-hurricane rebuild, the same data become a tool for public trust. And the World Bank's own record shows that returns depend less on the sector or the country than on how a project is designed, disbursed, and managed.

The dual hurdle: two tests for public capital

Henry began with a simple framework — the dual hurdle of Gardner and Henry (2023) — drawn as nested sets. Within the universe of investments a government might make sits the *socially desirable* subset: projects whose economic or social rate of return exceeds the government's opportunity cost of capital. Inside *that* sits the *commercially viable* subset: projects whose financial rate of return also clears a commercial cost of capital. Projects in the innermost circle satisfy both hurdles at once — good for society *and* attractive to private investors.

The point is allocation: if a socially valuable project can also earn a market return, private money should carry it, reserving public money for projects that deliver social returns but cannot attract private capital. To make those judgments in real time, the Hoover–World Bank Group "Learning from Lending" team is building two portals — one on financial rates of return, one on economic (social) rates of return. As Lall stressed, this is possible because nearly every World Bank investment figure is public, back to the first loan, and the economic-returns work is fully reproducible.

From archives to portals — and to Jamaica

Two portals

The financial-returns portal, building on Cole et al. (2024), turns decades of IFC equity investments into the metrics private investors actually use — multiples, internal rates of return, and public-market-equivalent hurdles benchmarked against the S&P 500 and the MSCI Emerging Markets index. Because the IFC can hold investments far longer than a private fund, it flags holding periods explicitly; the data remain confidential, so the demo uses dummy figures — itself an argument for transparency. The economic-returns portal draws on about 5,000 completed World Bank projects with rates of return from project-completion documents, sliced by region, sector, income group, and time.

The Jamaica use case

This is not an academic exercise. Henry chairs the Jamaica Reconstruction and Resilience Oversight Committee (JAMRROC), named by Prime Minister Andrew Holness to guide the rebuild after Hurricane Melissa with speed, foresight, and public trust. In the Jamaican data, the sectors with high financial returns — information and transportation in particular — also show high economic returns: where private investors can profitably build, public money is freed for schools and sanitation. In a setting where "profit" has long been a loaded word, the portals are as much a communication tool as an analytic one, letting leaders manage the whole *portfolio* so socially oriented early choices are balanced by higher-return projects later.

Evidence at a glance

- **Returns clear the bar — and rise over time.** About 85% of IBRD projects historically posted economic rates of return above the hurdle rate, and the median project shows only about a 3-point optimism bias (ex-ante social return about 21%, ex-post about 18%).
- **Direct benefits only.** World Bank economic rates of return count only directly attributable benefits — a transport project's travel-time and cost savings, not spillovers or land-price gains — so the estimates are a conservative lower bound.
- **Managers move the needle.** Over 60 years, IFC commitments are under about 5% of World Bank Group commitments — leaving wide scope to route the dual-hurdle test across institutions — and work by Lall and co-authors shows higher-quality managers raise economic returns, rising by roughly 12 points where supervision is strong.
- **Design beats sector.** Gill argued from the Bank's evidence that how a project is designed and implemented matters far more than which sector or country it is in — a well-designed project in a low-return sector can beat a poorly implemented one in a favorable sector. The framework also travels: Morocco has 286 projects in the database, 66 with economic rates of return.

From analysis to trust

Everyone is entitled to their own opinions, not their own facts. The dual hurdle is a decision rule for capital and a device for legitimacy — but its power depends on getting data quickly, managing the whole portfolio, and pairing the technology with political buy-in.

- **Route capital by the two hurdles, and think in portfolios.** Use private money for projects that clear both social and financial tests; reserve public and concessional money for socially valuable ones that cannot attract it, balancing the mix to avoid a debt-to-GDP problem.
- **Invest in the people who deliver.** Faster early disbursement and strong supervision — not rushing to closure — drive returns; skilled, experienced task managers are the key interface with client governments.
- **Pair technology with legitimacy.** Push for transparency on returns so governments and investors can judge what is investable; JAMRROC's authority is modeled on Jamaica's Economic Policy Oversight Committee, which helped cut debt from about 150% to 60% of GDP by solving problems "around a table rather than in the streets."

Day 8 • Wednesday, August 5, 2026 • Washington, D.C.

Official Launch of the World Development Report 2026: The Promise of Artificial Intelligence

A lecture by **Gaurav Nayyar**, Director, World Development Report 2026, World Bank Group.

The bottom line. The World Development Report 2026 reframes AI from a technology story into a development story. Its message to developing countries is threefold: be **optimistic**, because AI's capabilities can compress into a decade what once took a century; be **strategic**, by starting from adoption but treating *adaptation* to local data and conditions as the main task; and be **pragmatic**, making cost-conscious choices as enablers, users, and regulators of AI rather than chasing frontier systems or an "AI arms race."

The mindset: why developing countries should be optimistic

Most of what we read about AI concerns advanced economies; the report deliberately asks what AI means for development. It treats AI as a general-purpose technology — joining steam power, electricity, computers, and the internet — best understood through three features: the *capabilities* it brings, the *concentration* in its production, and the *complements* it needs to work in practice.

That optimism rests on capabilities that are scarce in many developing countries. Because AI can help solve complex problems, it "can do in a decade what might otherwise have taken a century." Early evidence is concrete: AI-generated weather forecasts in India help both optimistic and pessimistic farmers adjust their production and investment choices; Kenya's judiciary uses AI to allocate court cases to mediators better than the previous manual process; and in Bangladesh, AI-supported medical imaging is raising patient screening by almost a third to 40% per day amid a global shortage of radiologists. Most of these gains come through public-service delivery and often from older AI models — predictive analytics, remote sensing, early-warning systems — not chatbots; meanwhile job disruption is less imminent than in high-income countries, where a larger share of jobs is amenable to automation.

The strategy: adopt, but above all adapt

The good-news story has two caveats. Concentration is one: nearly all leading AI companies sit in just two countries — the United States and China — across every layer of the value chain. The report's answer is not to chase "AI sovereignty," which is too expensive for most developing countries and would only fragment the technology; it is to navigate the dependency with interoperable systems and to use others' massive investments to customize models locally. The bigger risk is a lack of complements: without the basic infrastructure, skills, and institutions, adoption — and therefore AI's productivity and growth benefits — will be slow to take off, which is why projected gains are larger in advanced economies at current adoption rates.

This is why adoption alone is not enough. AI is unique in that it learns from data, and data must reflect local languages, needs, and conditions; simply importing a tool from abroad may not work and can introduce bias. Advancing the frontier is relevant only for a small number of well-resourced

countries; for everyone else, adopting available AI is the logical start, but *adapting* it to context unlocks the value. This echoes the "investment–infusion–innovation" arc of the middle-income-trap report, with adaptation as infusion — except that with AI, adaptation must begin almost in parallel with adoption.

Evidence at a glance

- **Excitement tracks development needs.** Across about 30 countries, Asia and Africa are more excited than concerned about AI, while Europe and the Americas are more concerned than excited.
- **Capabilities in action.** India (AI weather forecasts for farmers), Kenya (AI case allocation in the judiciary), and Bangladesh (AI medical imaging raising screening by almost a third to 40% per day) show AI addressing long-standing development gaps.
- **Analog complements bind hardest.** In a firm survey across several developing countries, businesses of all sizes cite lack of knowledge, information, awareness, or finance — not just digital gaps — as what holds AI adoption back.
- **Pilots are not proof.** Of about 10,000 studies on AI in healthcare, fewer than 1% offered meaningful evidence that AI improved services — the world is full of experiments, not scaled solutions.
- **The least prepared may gain most.** About 800 voluntary AI standards are already published or in development; yet while all high-income and most middle-income countries have a national AI strategy, only 1 of 25 low-income countries does.

The policy agenda: enabler, user, regulator

Make it your own. The report reframes a technology question as a development-policy question. Governments have three roles to play — and the right sequence is to build foundations, learn what works, and earn trust, all under tight budgets.

- **Enable.** Provide the foundational infrastructure and skills — connectivity, electricity, local-language data, and some compute — recognizing that analog complements matter even more than digital ones.
- **Use.** As major buyers of AI, governments should reform procurement so it is continuous rather than one-and-done and not simply cheapest-bid, with frameworks to procure, evaluate, and scale the solutions that work.
- **Regulate.** With limited capacity and fast-moving technology, start with voluntary rather than mandatory standards, avoid rushing to legislate economy-wide, and make standard-setting more inclusive.
- **Favor open source, avoid lock-in.** Open models let countries truly customize; India, the UAE, and others have built local language models on open-source foundations that outperform frontier models in their own languages — but keep systems interoperable to avoid dependence on any one vendor. The report is the third in a trilogy — the middle-income trap, standards, and the promise of AI — built on one framework; next year's turns to jobs.

How Ready Are We for AI: New Insights from the BReady4AI Index

A lecture by **Ufuk Akcigit**, Co-Director, Growth Academy; Arnold C. Harberger Professor of Economics, University of Chicago (on leave); Deputy Chief Economist and Director of Private Markets, World Bank Group.

The bottom line. The headlines about AI come from large incumbents laying off workers, but most of the economy — and most of the developing world — is made of very small firms. New microdata show that the smallest businesses are the most likely to adopt AI and that adopters see more job creation than destruction and higher revenue productivity. And a new, living "BReady4AI" index reveals that countries' national AI strategies often do not match their real bottlenecks — making awareness of those gaps the first task of readiness.

Creative destruction meets a general-purpose technology

Akcigit opened where the Growth Academy began — with creative destruction. As with past general-purpose technologies, AI's cost-reducing effect will land first, displacing jobs, before the economy adjusts. Whether that adjustment is benign depends on dynamism — on new entry and young people bringing ideas we cannot yet imagine. If economies instead cling to old goods and simply make them more efficient, that is process innovation — and process innovation costs jobs. The optimism, then, rests on *product* innovation and on firms that scale.

The trouble is that almost everything we know about AI comes from large incumbents — Amazon, Oracle — yet firms with nine or fewer workers make up about 80% of U.S. businesses, and developing-country firms are smaller still. The left tail of the U.S. firm-size distribution is therefore the most informative window into AI's likely path in developing countries. With co-authors including James Evans and Enrico Colonnelli, Akcigit turned to real-time data from Intuit QuickBooks — the largest online platform for small businesses, with more than 10 million customers — complemented by the team's own surveys.

What the smallest firms reveal

Surveys say about 70% of small-business owners "use" AI — but usage is a fuzzy word. Following the logic that firms pay for what creates value (much as patent-renewal fees reveal a patent's worth), the team looked at spending: only about 10% of small businesses pay $10 a month and just under 1% pay $100 a month. Adoption is real but early and experimental. It is highest in information and professional services and lowest in agriculture — a sobering fact for developing economies where agriculture looms large.

The most striking pattern is a U-shape: adoption is *highest* among the very smallest firms. A one- or two-person business that cannot afford to hire — adding a worker can nearly double its costs — is the most desperate for help, and AI offers it cheaply. Ambition matters too: firms whose owners say they *want* to grow are far more likely to adopt than "subsistence" ones. And the outcomes echo creative destruction: most adopters report no change in employment, but among those that change, job creation exceeds destruction and revenue growth outpaces decline — so revenue productivity rises. The main barriers are privacy concerns and, tellingly, lack of knowledge, which correlates strongly with the entrepreneur's education.

Evidence at a glance

- **Small firms dominate.** Businesses with nine or fewer workers are about 80% of U.S. firms; the developing-world equivalent is smaller still, making the U.S. left tail the best available proxy.
- **Adoption is early.** About 70% of owners report using AI, yet only about 10% pay $10 a month and just under 1% pay $100 a month — firms are still experimenting.
- **The U-shape.** The likelihood of adopting AI is highest for one-worker firms, then falls with size — the smallest, most financially constrained businesses adopt most.
- **Ambition predicts adoption.** Transformative, growth-oriented entrepreneurs are much more likely to adopt AI than subsistence firms content to stay stable.
- **Creation beats destruction.** Among firms that change, AI adopters report more job creation than destruction and stronger revenue growth — a net gain in revenue productivity — while power outages and data-center gaps constrain adoption.

The BReady4AI Index — and its cautions

Every country is like a fingerprint. A living index — to be released on bereadyforai.com — lets policymakers track where nations stand as AI evolves; but its measures are indicators of *what countries do and say*, not verdicts on welfare or productivity.

- **Measure the whole ecosystem.** The index draws on academic papers, patents, GitHub activity, data centers, and government AI grant registries — with syllabi and legal documents being added — across four pillars: frontier creation, global influence, diversity and collaboration, and compute and talent. Early scorecards rank the U.S. first, China second, and the U.K. third.
- **Read national strategies in three dimensions.** Processing 80-plus national AI strategies by frequency, intensity, and direction (expansionary versus safeguarding) shows sharp contrasts — the U.S. strategy is dominated by security and safeguarding, China emphasizes R&D and industrial adoption expansively, and the EU leans toward safeguarding regulation.
- **Watch for mismatch.** Combined with the World Bank's governance indicators, the data show weaker rule of law linked to more emphasis on controlling AI, and — the red flag — some countries with frequent power outages whose strategies barely mention energy.
- **Do not read expansion as "good."** As Akcigit and Evans cautioned, expansionary effort can breed market concentration that kills competition and future innovation, and stated strategies may diverge from actual laws — so the index is a starting point for analysis, not a scorecard of success.
- **Build the human-capital complement next.** The planned education pillar reflects the core message: AI pays off when combined with other skills, so what matters is whether farmers and health workers — not just AI specialists — can adopt it.

Learning from 80 Years of Lending: Economic Returns and the People Behind World Bank Projects

A lecture by **Somik Lall**, Co-Director, Growth Academy; Director, Strategy and WBG Institute for Economic Development, Development Economics Vice Presidency, World Bank Group.

The bottom line. The best way to know what an institution does is to see where it puts its money. Eighty years of World Bank infrastructure lending have delivered high economic returns — about 24% on average — and, surprisingly, the returns are *highest* in the poorest-governed, most capital-scarce countries. What makes the difference is not the country's institutions alone but the World Bank's own performance, and above all the individual project leaders who make investments deliver on the ground. The uncomfortable corollary: the institution rewards those people far too little.

Put the money where the mouth is

You can learn from research and from practitioners, but the surest test of what an institution believes is where it actually lends. The World Bank's first loan was a one-page letter from the government of France on 25 June 1946 — $500 million for reconstruction, with no economic analysis and no rationale, just a request for help. Eighty years later, the question is whether all that lending has paid off. To answer it, Lall and colleagues in the research and development-economics group synthesized the social (economic) rates of return on IBRD- and IDA-funded infrastructure over time and across sectors.

This is deliberately different from the work presented the day before by Peter Henry, which benchmarked the IFC's *financial* return against yardsticks like the S&P 500. The IFC is an impact investor, not a competitor to standard investors; the IBRD's mandate is economic and social value, not private-market profit. So this exercise measures economic value, and every input is public: the underlying investment data can be downloaded and re-run through a reproducibility package. As Lall put it, "your country paid for this — this is your data." It matters because infrastructure is a long-lived bet: a road or dam delivers for 60–70 years, a very different frame from a health or education project whose benefits show up over 8–10.

Systems, and the people who beat them

The standard aid-effectiveness literature says aid works only where governance, rule of law, and institutions are strong — which implies that fragile and conflict-affected economies should earn low returns. Lall calls that nonsensical, because those are exactly the places that need concessional finance most. What makes high returns possible in hard places is the *system* an institution puts around a project — and, crucially, the individuals who run it. Much of the measured return cannot be attributed to systems at all; echoing John List, it depends on people who are hard to scale: the task team leaders who work inside project management units, manage procurement, and make sure the money builds the road rather than ending up in a Swiss bank account.

To see this, the team built an unusually deep dataset: economic rates of return for completed projects over 60 years (available for about half of them), matched for the first time to loan-disbursement records from borrowers' own treasuries, to 50 years of confidential CPIA governance scores, and to the scraped career histories of 1,600 individual task team leaders drawn from archives reaching back to the 1940s.

Evidence at a glance

- **Returns are high and rising.** Economic rates of return average about 24% (median 18.5%); after a structural break around 1997 they rose from roughly 15% to 21% — the opposite of the usual finding that infrastructure returns decline over time.

- **Poorly-governed countries do better.** Countries below the median CPIA earn *higher* returns than better-governed ones, suggesting scarce capital is flowing where needs are greatest.
- **Institutional performance pays.** In weak-governance countries (CPIA below three) base returns run about 23%; strong World Bank supervision and stronger project units add roughly 12 percentage points.
- **Timing is asymmetric.** Getting the first disbursements out early raises returns, but rushing a project to close does not — the final stages are where monitoring catches the roads and bridges that were never built.
- **The leaders are exceptional.** About 30% of task team leaders hold PhDs; only 5% are hired locally; 10% come through the Young Professional Program (about 100 chosen from 200,000 applicants). A typical infrastructure project runs about $90m over roughly 2,600 days (seven years) at around 3.5% interest.

What it means for the institution

Value comes from capability, not just capital. The Bank adds value less by moving money than by building the supervision, accountability, and teams that turn a loan into delivered returns — and it should reward the people who do that work.

- **Fix the talent gap.** The Bank practices negative assortative matching — sending its best people to its hardest environments, where they deliver returns near 40% — yet those same experts average about 1.5 promotions against a possible six. Create tracks that let them keep technical expertise while gaining real influence.
- **Do not confuse speed with success.** Pressure to compress implementation timelines is misguided; development is a long business, and cutting supervision short is where projects fail.
- **Treat the estimates as conservative.** Returns count only directly attributable social value (for example, reduced travel time and trade costs), excluding second-order effects such as land prices — so the true benefits are likely larger.
- **Keep a dual focus.** Maximize economic value through better project design in high-need contexts *and* pursue broader institutional reform; the two are complements, not rivals. Getting results on the ground is hard, and the participants' own task is to build a chain from problem to action to scale.