inference.

The GAIN Framework

Gross AI-Native capacity of Nations

and the National AI Accounts (NAIA)

A national operating system for the AI economy

Measuring, governing, and accelerating national wealth creation in the age of artificial intelligence. Developed as a next-generation successor to first-wave AI measurement systems, including the Stanford Digital Economy Lab's AI Economic Indicators.

INFERENCE RESEARCH · INAUGURAL EDITION · JUNE 2026

Know, don't guess.

Contents

Preface: Why Nations Need AI Accounts

In 1937, Simon Kuznets delivered to the United States Congress the first national income accounts. The Great Depression had been raging for eight years, yet until that moment no government on earth could say, with any precision, how large its economy was, how fast it was shrinking, or whether its policies were working. The instruments came after the storm had already broken. The System of National Accounts that grew from Kuznets's work became the most consequential measurement technology of the twentieth century: it did not merely describe economies, it made modern economic management possible. What gets measured gets governed.

Artificial intelligence now presents an equivalent measurement emergency, but with the sequence reversed: this time the storm is visible while it forms. The Stanford Digital Economy Lab's AI Economic Indicators project — whose June 2026 inaugural research note documents early-career employment contracting at 3.8% annually in AI-exposed occupations while a "Takeoff Tracker" finds the capital share of income grinding steadily upward — represents the Kuznets moment of the AI era: the first serious attempt to build measurement infrastructure before the transformation completes rather than after. Erik Brynjolfsson's warning in that note is exact: "*We are flying blind into one of the most consequential periods in world history.*"

But the Stanford indicators, like all first instruments, measure what is easiest to see — employment, adoption surveys, aggregate productivity series — and they measure it for one country, descriptively, without a theory of how a nation converts AI into durable wealth. They are a barometer. Nations need a navigation system.

This report introduces that system: the **GAIN Framework — Gross AI-Native capacity of Nations** — and its operational instrument, the **National AI Accounts (NAIA)**. The ambition is deliberate and stated plainly: to do for the AI economy what the System of National Accounts did for the industrial economy. GDP measured the flow of industrial output; GAIN measures a nation's stock of AI productive capacity and the returns it realizes from that capacity. The framework comprises seven capital accounts, a headline composite score, a country classification typology, a value-creation model with identified intervention points, a national flywheel diagnostic, and a decade-length policy playbook — applied throughout to ten nations, with India developed as a full case study.

The report is organized in nine parts, following a single arc: what existing instruments miss (Part 1), what the global landscape actually looks like (Part 2), what a complete measurement system requires (Parts 3–4), how AI becomes national wealth (Parts 5–7), and what governments should do and measure (Parts 8–9).

A note on epistemics. All country assessments are inference-based syntheses of public evidence as of June 2026; scores presented are illustrative calibrations intended to demonstrate the methodology, not audited statistics. The framework is designed precisely so that, once adopted, such numbers would cease to be estimates and become accounts.

Part 1. Critical Review: The First Instruments and Their Blind Spots

1.1 What problem is Stanford trying to solve?

The Stanford Digital Economy Lab's AI Economic Indicators (AIEI) project attacks a specific institutional failure: the latency, granularity, and conceptual mismatch between official statistics and the speed of AI-driven change. Official labor statistics arrive quarterly with months of lag, at occupational granularity too coarse to detect technology-specific displacement; productivity statistics are annual, heavily revised, and notoriously silent during the early diffusion phase of general-purpose technologies. AI capability, by contrast, compounds on a cadence of months. The Lab's stated purpose is "to make AI's economic effects broadly visible" — to replace anecdote with high-frequency, occupation-level, age-disaggregated evidence so that policymakers, firms, and workers can respond to the transformation rather than discover it retrospectively.

This is, in effect, an attempt to build an early-warning system. The naming of its flagship product — the *Canaries Dashboard* — is candid about the epistemology: the project looks for the most sensitive members of the labor market (early-career workers in AI-exposed occupations) on the theory that they will register disruption before aggregates do. It is measurement as sentinel surveillance, borrowed conceptually from epidemiology.

1.2 What is currently being measured

The June 2026 inaugural research note assembles three instruments:

Instrument	What it measures	Headline findings (June 2026)
Canaries Dashboard (with ADP Research)	Monthly employment trends across ~730 occupations in a balanced panel of 25,000 firms (4.6M matched workers), cross-tabulated by AI exposure quintile (Eloundou et al. 2024) and worker age; usage character via the Anthropic Economic Index automation/augmentation ratios	Aggregate divergence modest (most-exposed +1.1%/yr vs least-exposed +2.0%/yr); early-career (22–25) workers in exposed occupations contracting at −3.8%/yr; automation-skewed usage correlates with decline, augmentation usage does not
Takeoff Tracker	12 aggregate U.S. indicators of "explosive growth" — capital share, TFP growth, IP-equipment share of capital stock, and related macro series	7 of 12 indicators show no evidence of takeoff; 3 mild; 2 strong (capital share rising persistently; IP-equipment share recovering to early-2000s levels); no decisive takeoff
Adoption Monitor	Survey-based individual and firm adoption (Bick–Blandin–Deming, Gallup, Pew, Yotzov et al. four-country firm survey)	Individual adoption trending up with a recent workplace reversal; firm adoption widespread and U.S.-led; large current-vs-expected gaps in robotics and autonomous vehicles

The intellectual architecture is recognizably Brynjolfssonian: the automation/augmentation distinction operationalizes his "Turing Trap" thesis (2022) — that economies which use AI to *imitate* human labor reap displacement, while those that use it to *complement* labor reap broadly shared gains. The Canaries finding that only automation-skewed usage correlates with employment decline is the first dashboard-grade evidence for that thesis.

1.3 A researcher-grade critique: seven blind spots

The project deserves to be judged by the standard it sets for itself — becoming "the most reliable, timely source for understanding AI's economic impact." Against that standard, seven structural gaps stand out. These are not defects of execution; most reflect deliberate scoping. But together they define the space a successor framework must fill.

Blind spot 1: It measures weather, not climate — effects, not capacity. Every AIEI instrument observes *outcomes* of AI diffusion (employment shifts, adoption rates, macro ratios). None measures the *productive capacity* that generates those outcomes: compute stock, talent depth, data assets, model capability owned domestically. A nation could score identically on every Stanford indicator while owning none of the underlying means of AI production — and the dashboard would not notice. This is the difference between a thermometer and a balance sheet.

Blind spot 2: One-country instrumentation for a planetary phenomenon. The Canaries panel is U.S. ADP payroll data; the Takeoff Tracker uses U.S. macro series; only one adoption survey spans four countries. Yet the defining economic fact of AI in 2026 is its *international* structure: $285.9 billion of U.S. private AI investment versus $12.4 billion in China (with an estimated $184 billion in Chinese state guidance funds behind the official figure); Gulf states committing tens of billions to gigawatt-scale compute campuses; an 89% collapse since 2017 in AI researcher migration to the U.S. There is no instrument for comparing nations, and therefore no instrument for the competition that most preoccupies the governments the Lab hopes to inform.

Blind spot 3: The balanced-panel design filters out the most important margin. By construction, the Canaries sample excludes firm entry, exit, and switching — precisely the channels through which general-purpose technologies historically reallocate economic activity. The AI-native startup founded in 2024 is invisible; so is the incumbent that fails because of one. The dashboard is well designed to detect *within-firm* labor substitution and structurally unable to detect *Schumpeterian* creative destruction.

Blind spot 4: Exposure is not usage, and usage is not value. The exposure scores (Eloundou et al.) are model-derived potentials, not observed behavior; the Anthropic Economic Index covers one vendor's traffic; and as the Lab itself concedes, current exposure measures "do not account for how demand responds to price changes induced by AI-driven productivity gains." A measured employment decline in an exposed occupation could reflect displacement, demand elasticity, offshoring, or interest rates.

The causal chain from AI capability to economic outcome runs through task-level cost changes — and no link of that chain is directly priced anywhere in the system.

Blind spot 5: No fiscal, distributional, or geographic ledger. Who captures AI's returns — capital or labor, founders or workers, San Francisco or the median county, the treasury or offshore entities — is the political economy question of the decade. The Takeoff Tracker's rising capital share gestures at it; nothing in the system decomposes it.

Blind spot 6: No prices, no trade, no external account. AI inputs and outputs now trade internationally at observable prices: GPU-hours, API tokens, model licenses, AI-embodied services. A nation can run a surplus or deficit in artificial intelligence exactly as it can in oil. No existing indicator system — Stanford's or anyone's — constructs an AI current account or AI terms of trade. The single most consequential macroeconomic variable of the 2030s is unmeasured.

Blind spot 7: Description without a steering model. The AIEI tells a policymaker *what is happening* but embeds no theory of *what to do*: no identified intervention points, no leading-versus-lagging structure, no notion of which national capacities cause which outcomes with what time constants. It is — by design — a research observatory, not an operating system.

1.4 Why GDP-era metrics fail in the AI era

The deeper problem precedes Stanford: the canonical quartet of GDP, labor productivity, employment, and patent-based innovation metrics was engineered for an economy of *priced, counted, human-produced* goods. AI violates each design assumption in turn.

GDP mismeasures free and near-zero-marginal-cost output. GDP records transactions at market prices. When an AI system delivers at near-zero price what previously cost hundreds of dollars — translation, legal drafting, tutoring, diagnosis — measured GDP can *fall* while welfare rises. Brynjolfsson's GDP-B research program quantifies the wedge: by early 2026, estimated U.S. consumer surplus from generative AI alone reached roughly $172 billion annually (up from $112 billion a year earlier), nearly all of it invisible to national accounts. The industrial-era identity "output ≈ welfare ≈ income" is breaking apart, and each fragment now needs its own instrument.

Productivity statistics suffer the GPT J-curve. As Brynjolfsson, Rock, and Syverson showed, general-purpose technologies require massive complementary intangible investment — process redesign, data infrastructure, skills — that national accounts book as *cost* rather than capital formation. Measured TFP therefore stagnates or dips precisely while transformative capacity accumulates, then surges years later. The Takeoff Tracker's "neutral" TFP reading in 2026 is exactly what the J-curve predicts at this phase — which means TFP is uninformative as a contemporaneous signal, and capacity must be measured directly.

Employment counts miss composition, quality, and the broken career ladder. Aggregate employment can hold steady while the *entry rungs* of professional careers are removed — the precise pattern the Canaries data documents (−3.8%/yr for exposed early-career occupations against +2.0% for unexposed). A headline unemployment rate is structurally blind to a generation locked out of skill-building roles, to within-occupation task hollowing, and to wage-distribution shifts between AI-leveraged and AI-displaced workers.

Patent and R&D metrics misread AI innovation. Frontier AI advances are disclosed in preprints and open weights, not patents; the binding inputs are compute and talent, not booked R&D. China leads the world in AI patent volume while the frontier of capability is contested between a handful of U.S. labs and Chinese open-weight challengers whose key artifacts are deliberately unpatented. Counting patents in 2026 is like counting monasteries to measure literacy after the printing press.

1.5 The invisible economy: what national accounts cannot see

Drawing the threads together, at least eight categories of AI-generated value are currently invisible or mispriced in national accounts: consumer surplus from free AI services; intangible complementary capital (workflow redesign, fine-tuned models, prompt libraries, agent infrastructures booked as expenses); data assets, which appear nowhere on any national balance sheet despite being a primary factor of AI production; quality improvements delivered at constant price (the same doctor's visit, now AI-checked); time savings in unpaid household and informal-sector production; open-source model value, where a freely downloaded frontier-class model transfers billions in capability at a recorded price of zero; depreciation asymmetries (AI capital depreciates in 2–3 years while statistical conventions assume far longer, overstating net capital); and option value — the national strategic value of *possessing* frontier capability, which behaves like a real option and is priced by markets in sovereign-adjacent equity valuations but by statisticians not at all.

The conclusion of this review is not that the Stanford indicators are wrong — they are the best first instruments in existence, and this report builds gratefully on them. The conclusion is that they are *partial by design*: a single-country observatory of effects, when what governments now require is a multi-country accounting system of capacities, flows, and returns, joined to a causal model that tells them where to act. That is the gap the remainder of this report fills.

Part 2. The Global AI Economy Landscape

2.1 The shape of the race in mid-2026

Three facts organize everything else. First, *capital concentration*: U.S. private AI investment reached $285.9 billion in 2025 — roughly 23 times China's reported $12.4 billion, though Chinese state guidance funds (an estimated $184 billion deployed since 2000) make the true gap far smaller than private figures suggest. Second, *capability convergence*: the U.S.–China frontier-model performance gap has narrowed to single digits (roughly 2.7 percentage points on composite benchmarks as of March 2026), with Chinese open-weight models diffusing globally at zero price. Third, *the rise of the compute petrostates*: the UAE's Stargate campus (5 GW planned, first gigawatt costing $30+ billion) and Saudi Arabia's HUMAIN (1.9 GW by 2030, $23 billion in supplier deals) represent a new strategy — converting hydrocarbon wealth and sovereign energy into AI infrastructure rents.

Around these poles, middle powers are executing four distinguishable strategies, profiled below across the ten assessment dimensions: compute infrastructure, talent density, startups, investment, adoption, digital public infrastructure (DPI), government policy, AI exports, data assets, and overall readiness.

2.2 Summary assessment matrix

Ratings are inference-based on a five-point scale (5 = global frontier, 1 = nascent), synthesizing the AI Index 2026, OECD.AI policy trackers, national budget documents, and investment data as of June 2026.

Country	Comp.	Talent	Startup	Invest	Adopt	DPI	Policy	Export	Data	Ready
United States	5	5	5	5	5	2	3	5	5	5
China	4	5	4	4	5	4	5	4	5	5
UAE	4	2	2	4	4	4	5	3	2	4
Singapore	3	4	3	3	5	5	5	3	3	5
South Korea	4	4	3	4	4	4	5	4	3	4
United Kingdom	3	4	4	4	4	3	4	4	4	4
Saudi Arabia	3	2	2	4	3	3	4	2	2	3
Canada	3	4	3	3	3	2	4	3	3	3
France	3	4	3	3	3	3	4	3	3	4
India	2	4	3	2	3	5	4	4	4	3

2.3 Country profiles

United States — the full-stack producer

The U.S. owns every layer of the AI stack at frontier scale: the chip design monopoly (NVIDIA), the leading cloud hyperscalers, the majority of frontier labs, the deepest capital markets ($285.9B private investment; 1,953 newly funded AI companies in 2025, ten times the next country), and the largest measured consumer surplus (~$172B/yr). Adoption is the global benchmark — U.S. firms lead every category in the four-country Yotzov et al. survey cited by Stanford. Its vulnerabilities are equally structural: energy and grid constraints on data-center buildout; an 89% collapse in net AI-talent immigration since 2017 (an 80% decline in the last year alone), eroding its historic talent-magnet advantage; near-absent digital public infrastructure (no national digital ID, fragmented payments); federal policy oriented toward export control and infrastructure permitting rather than diffusion support; and extreme geographic and firm-level concentration of gains. The U.S. is the archetypal **AI Producer** — it manufactures the means of AI production for the world and prices them.

China — the parallel-stack producer

China is running a deliberately different race: behind at the frontier of training compute (constrained by export controls, compensated by Huawei Ascend 950 clusters, Cambricon's 160% revenue growth, and conditional access to NVIDIA H20s), but ahead on diffusion, open-weight strategy, and physical-world AI. DeepSeek's R1-to-V4 trajectory demonstrated near-frontier capability at a fraction of Western cost, and the state's "AI+" initiative explicitly subsidizes adoption across manufacturing, logistics, and government — where China already leads the world in industrial robot installations. It leads in publications, citations, and patent volume; its data assets are vast and state-accessible; its DPI (payments, identity, super-apps) is world-class. Constraints: leading-edge semiconductor manufacturing, capital-market depth, and a talent system that now retains domestically but attracts few foreigners. China is simultaneously an **AI Producer** (second full stack) and the world's most aggressive **diffusion state**, exporting its stack to the Global South at concessional prices.

United Arab Emirates — the sovereign platform

The UAE has executed the purest version of the platform strategy: leveraging energy, capital, land, and geopolitical positioning to become an indispensable *host* of AI capacity it did not invent. Stargate UAE (G42, OpenAI, Oracle, NVIDIA, SoftBank; 1 GW initial, 5 GW campus, $30B+ first phase, 200 MW live 2026–27), the Falcon open-weight model family, MBZUAI (the world's first AI-dedicated university), a national AI Minister since 2017, and AI-native government services. Weaknesses are the mirror of its strengths: a small indigenous talent base (~10M population, heavily expatriate), thin startup depth, limited proprietary data assets, and strategic dependence on U.S. chip-access diplomacy. The UAE is the archetypal **AI Platform**: it monetizes location, energy, and capital as compute rents and aims to convert them into ecosystem gravity.

Singapore — the orchestrator state

Singapore tops AI-readiness rankings (≈0.80 readiness score) by treating AI as a whole-of-government operating discipline: NAIS 2.0 strategy, AI Verify governance tooling exported as a global standard, SEA-LION regional language models, sovereign GovTech stack, and deliberate rationing of land, power, and carbon for data centers (capacity ~1.4 GW, deliberately constrained). Talent density per capita is exceptional; adoption in finance, logistics, and government is world-leading. Its limits are physical: energy and land cap its compute ambitions, its domestic market is tiny, and it produces few frontier artifacts. Singapore is the archetypal **AI Ecosystem Builder** — a small state that competes on institutional quality, trust, and orchestration of other people's capabilities, exporting governance itself as IP.

South Korea — the embodied-AI industrialist

Korea pairs the world's densest AI patenting per capita with the industrial base to put AI into objects: memory-chip dominance (Samsung, SK Hynix supply the HBM that every frontier GPU requires), the world's highest industrial-robot density, and national champions (Naver, LG, SK Telecom, Upstage) selected in 2025–26 to build sovereign foundation models under a $5.7B National Growth Fund allocation and a February 2026 AI Action Plan containing 99 execution tasks. GPU stock is being built toward the tens of thousands under direct state procurement. Constraints: language-market scale, startup capital depth, demographic decline, and dependence on U.S. model ecosystems at the frontier. Korea is a **Producer–Builder hybrid**: a component superpower converting itself into a sovereign-model state.

United Kingdom — the science-and-services producer

The UK holds Europe's deepest AI research base (DeepMind's legacy, Oxford/Cambridge/Edinburgh), its largest startup ecosystem (~885 funded AI startups, third globally), the AI Safety/Security Institute as a governance export, and roughly $4.5B in committed public AI investment plus a compute buildout (Isambard-AI). Its December 2025 accession to the Pax Silica framework formalized infrastructure access in exchange for alignment with U.S. export architecture. Weaknesses: chronic scale-up capital flight to the U.S., grid constraints, no hyperscaler of its own, and a productivity-diffusion problem in its SME long tail. The UK is a **Producer of science and standards** that persistently struggles to retain the wealth its science creates — the canonical case of strong innovation capital with weak value capture.

Saudi Arabia — the capital-scale platform

Saudi Arabia is executing the UAE playbook at larger capital scale and earlier institutional maturity: HUMAIN (PIF-backed; up to 1.9 GW by 2030; $23B in deals with NVIDIA, AMD, Qualcomm, AWS; a $3B AirTrunk data-center partnership), SDAIA as central AI authority, Arabic LLMs (ALLaM), and AI provisions woven through Vision 2030 giga-projects. Its advantages are energy (the world's cheapest at scale), sovereign capital, and demographic youth. Its gaps: talent density, research depth, private startup formation, data openness, and institutional trust ratings. Saudi Arabia is an **AI**

Platform one stage earlier than the UAE — currently buying the stack, with the open question being whether rented capability becomes owned capability.

Canada — the inventor who rents out the workshop

Canada invented much of deep learning (Hinton, Bengio; CIFAR's pan-Canadian strategy was the world's first national AI strategy, 2017) and retains elite research institutes (Vector, Mila, Amii) and a genuine frontier-adjacent champion in Cohere ($240M anchor government investment). The 2024–26 Sovereign AI Compute Strategy (CAD $2B: $700M private compute challenge, ~$890M–$1B public supercomputer beginning FY2026-27, $300M access fund) is a direct response to its core failure mode: decades of inventing capabilities that scaled into American firms. Adoption by Canadian industry remains mediocre; energy abundance is an underexploited platform asset. Canada is an **Ecosystem Builder by necessity** — research-rich, diffusion-poor, fighting talent gravity from a neighbor ten times its size.

France — the sovereign producer of Europe

France has pivoted hard from regulatory leadership to industrial policy: Mistral AI as Europe's only credible frontier-model company, the February 2025 AI Action Summit's €109B mobilization of (largely Gulf and private) capital for French compute, cheap nuclear baseload as a decisive energy advantage, and elite mathematical talent pipelines (École Polytechnique, ENS) that staff frontier labs worldwide. Constraints: EU AI Act compliance burden it must carry while competing against lighter-touch regimes, fragmented European capital markets, and slow enterprise adoption. France is Europe's **Producer candidate**: the only EU state plausibly fielding a full sovereign stack — nuclear energy, frontier lab, national cloud — by 2030.

India — the diffusion superpower in waiting

India enters the AI era with the world's most distinctive asset mix: the deepest AI-skilled labor pool outside the U.S. and China (millions of developers; the largest national cohort of GitHub AI contributors), the world's most advanced population-scale DPI (Aadhaar identity ~1.4B, UPI processing on the order of 19–20 billion transactions monthly, DEPA data-sharing architecture), a $1.25B IndiaAI Mission that has taken common compute from zero to ~38,000–58,000 GPUs at subsidized rates (₹65–150/hour) with a 100,000-GPU target, and an IT-services sector that already exports AI-embodied services at scale. Its deficits are equally clear: compute per capita among the lowest of any major AI nation, modest private AI investment, no frontier model, research depth thin relative to talent volume, and chronic gaps between policy announcement and state-capacity execution. India is profiled fully in Part 6; in the typology it is the world's largest **AI Consumer** currently attempting the jump to **Ecosystem Builder**, with a plausible path to Producer status in narrow domains by the 2030s.

2.4 A classification model: the PPCE typology

Existing rankings collapse nations onto a single axis ("readiness"), obscuring the fact that countries are playing *different games*. The GAIN landscape model classifies nations on two structural axes:

- **Axis 1 — Capability Creation (CCI):** does the nation produce the means of AI production — frontier models, chips, training infrastructure, fundamental research?
- **Axis 2 — Absorption & Diffusion (ADI):** does the nation convert AI — its own or others' — into economy-wide productivity, public-service quality, and new firms?

Crossing the axes yields four archetypes, with a fifth strategic overlay:

Archetype	Definition	Economic model	2026 members
AI Producers	High capability creation; own frontier assets (models, chips, labs)	Sell the means of production; capture scarcity rents on compute, models, talent	United States, China; UK and France as partial/specialized producers; Korea at the component layer
AI Platforms	Import capability, host it at scale via energy, capital, and location	Earn infrastructure rents; convert hosting into ecosystem gravity over time	UAE, Saudi Arabia
AI Ecosystem Builders	Moderate creation, high orchestration: institutions, governance, DPI, talent programs deliberately compound	Compete on trust, standards, and absorption efficiency; export governance and niche capability	Singapore, Canada, Korea, France; India ascending
AI Consumers	Low creation; adopt imported AI with limited strategic accumulation	Take prices set elsewhere; gains accrue mostly to foreign capability owners	Most of the world by default; India is the largest consumer actively engineering its exit from this cell

The fifth element is **trajectory**: archetypes are positions, not destinies. The UAE is attempting Platform→Producer (Falcon, MBZUAI); Korea, component-Producer→full Producer; India, Consumer→Ecosystem Builder; the UK risks the reverse slide, Producer→Consumer-of-its-own-inventions, if value capture continues to leak. The classification's policy use is exactly this: a nation must first know which game it is in before it can know which pillar investments (Part 3) have the highest marginal return — and the scorecard in Part 4 is designed to make both position and trajectory measurable annually.

Part 3. The GAIN Framework: Gross AI-Native Capacity of Nations

3.1 The core idea

GDP answers the industrial era's central question: *how much does this nation produce?* The AI era's central question is different: *how much can this nation produce, decide, and invent per unit of human effort — and who owns the capacity that makes it possible?* Answering it requires measuring stocks as well as flows, capacity as well as output, and ownership as well as activity.

GAIN — Gross AI-Native capacity of Nations — is a national accounting framework built on three design principles:

1. **Account, don't rank.** Rankings flatter and obscure; accounts decompose and explain. GAIN is structured as a set of seven capital accounts — the **National AI Accounts (NAIA)** — from which a headline composite (the GAIN Score, Part 4) is derived, exactly as GDP is derived from the System of National Accounts rather than estimated directly.
2. **Separate capacity from returns.** Six accounts measure the *means* of AI production a nation possesses (capacity); the seventh measures what it actually *harvests* (returns). The gap between the two — the **Conversion Ratio** — is the single most diagnostic number in the framework: it distinguishes nations that own AI assets from nations that monetize them.
3. **Measure the border.** AI value crosses borders as chips, tokens, talent, models, and AI-embodied services. GAIN therefore includes the framework's most novel constructs: the **AI Current Account** (net trade in AI inputs and outputs) and the **AI Terms of Trade** (the price ratio of the AI a nation sells to the AI it buys).

The production-function logic underneath is explicit. National AI output capacity is modeled as:

$$Y(AI) = A \cdot [C^\alpha \cdot H^\beta \cdot D^\gamma] \cdot \Phi(\text{diffusion}) \cdot \Psi(\text{institutions})$$

where C is compute capital, H is cognitive (human) capital, D is data capital, A is innovation capability (the technology shifter), Φ is a diffusion multiplier capturing how widely capacity penetrates the economy, and Ψ is an institutional dampener/amplifier capturing trust, governance, and energy/regulatory throughput. The three bracketed factors are complements, not substitutes — which is why the scorecard aggregates them geometrically (Part 4): abundant compute cannot rescue absent talent, and vice versa. The seventh account, Realized Returns, measures Y itself and its distribution.

The seven accounts, in causal order:

#	Account	Measures	Account type
P1	Compute Capital	Sovereign-relevant stock and flow of AI computation and the energy	Physical capital

#	Account	Measures	Account type
		beneath it	
P2	Cognitive Capital	Depth, density, and net flow of AI-capable human talent	Human capital
P3	Data & Infrastructure Capital	National data assets and the digital public rails that make them productive	Intangible/public capital
P4	Innovation Capital	Capability to create new AI artifacts: research, models, startups, IP	Knowledge capital
P5	Diffusion Capital	Speed and depth with which AI penetrates firms, government, and households	Absorption capacity
P6	Institutional Capital	Trust, governance quality, regulatory throughput, societal license	Institutional capital
P7	Realized Returns	Productivity, surplus, fiscal yield, exports, and distribution actually harvested	Outcome account

3.2 Pillar P1 — Compute Capital

Definition. The stock of AI-relevant computation a nation can deploy on demand — owned, hosted, or contractually secured — together with the energy, grid, and semiconductor supply relationships that sustain and grow it.

Why it matters. Compute is the AI era's primary capital good: every model trained, every token served, every agent run is a draw against it. It is also the most price-volatile and most geopolitically gated input — allocated in 2026 as much by export-control diplomacy (Pax Silica, H20 licensing) as by markets.

Economic logic. Compute behaves like industrial-era electrification capacity: a general-purpose capital stock whose marginal product is realized only through complements (P2–P5), but whose absence is an absolute binding constraint. Because frontier compute depreciates in 2–3 years, the *flow* of investment matters as much as the stock — a nation that pauses buildout for 30 months effectively exits the frontier. Sovereignty matters because rented compute is subject to price discrimination and political recall: ownership structure enters the account, not just capacity.

Measurement methodology. Annual census of national AI compute, decomposed by: installed accelerator stock in FLOP/s-equivalents (training-class vs inference-class); effective availability (utilization-adjusted); ownership tier (sovereign-owned / domestically owned private / foreign-owned domestic / contractually secured offshore); energy headroom (MW of grid capacity permitted and queued for AI loads); and price (national GPU-hour reference price vs world benchmark). Normalize per $bn GDP and per capita for cross-country comparison.

Core indicators. Installed AI compute stock (FLOP/s-eq, absolute and per capita); compute self-sufficiency ratio (domestic capacity ÷ domestic demand); training-class share; national GPU-hour price index and its ratio to U.S. reference; data-center energy pipeline (GW permitted, under construction, operating); accelerator import dependence (HHI of supplier countries); sovereign compute access rate (share of researchers/startups reporting compute as non-binding).

Data sources. National energy regulators and grid queues; customs data on HS-coded accelerator imports; hyperscaler and colocation disclosures; Epoch AI and TOP500-style registries; cloud spot-price scrapes; national compute-mission administrative data (e.g., IndiaAI portal, EuroHPC).

Policy implications. This account tells a government whether to build, buy, or broker. Low stock + high energy headroom → platform strategy (build and host). Low stock + low energy → broker (secure offshore contracts, negotiate access treaties). High stock + low utilization → the constraint is P2/P5, not silicon — stop buying GPUs.

3.3 Pillar P2 — Cognitive Capital

Definition. The stock and net flow of humans who can build, deploy, govern, and economically apply AI — from frontier researchers through ML engineers to AI-fluent managers, clinicians, and civil servants.

Why it matters. Talent is the only AI input that is simultaneously a factor of production, a diffusion vector, and a compounding asset. The AI Index's finding that researcher migration to the U.S. has fallen 89% since 2017 signals the largest reallocation of strategic human capital since the post-war brain drains — and most nations cannot see their own position in it.

Economic logic. Cognitive capital exhibits increasing returns through agglomeration (researchers are more productive near other researchers) and through teaching externalities (each frontier engineer trains successors). It depreciates through skill obsolescence and emigrates at market speed. Critically, the Canaries evidence shows AI is eroding the *apprenticeship rung* of cognitive-capital formation — early-career roles — meaning nations must now account for whether their talent pipeline's entry stage is being automated away beneath them.

Measurement methodology. A talent stock-and-flow table, like a demographic account: stock by tier (T1 frontier researchers; T2 ML engineers/builders; T3 AI-fluent professionals; T4 AI-literate workforce), inflow (graduations, immigration, reskilling completions with verified competency), outflow (emigration, exit from field), and quality adjustment via publication/repository/credential signals. Density expressed per 10,000 workers.

Core indicators. AI talent stock by tier and density; net talent balance (annualized in/out migration of T1–T2); apprenticeship-rung index (early-career hiring share in AI-exposed occupations — directly importing Stanford's Canaries construct as a national indicator); top-venue publication authorship per million population; share of workforce with verified AI credentials; brain-circulation ratio (returnees ÷ leavers).

Data sources. LinkedIn/GitHub talent-flow data; OECD migration statistics; conference authorship registries; national credentialing bodies; payroll-data partnerships (the ADP model, replicated with national providers); university administrative data.

Policy implications. Distinguishes production problems (too few graduates → education investment) from retention problems (graduates leave → compensation, compute access, and agglomeration policy) from utilization problems (talent present but in low-AI roles → diffusion policy). Each has a different ministry and a different time constant.

3.4 Pillar P3 — Data & Infrastructure Capital

Definition. The nation's economically usable data assets — public, private, and personal — together with the digital public infrastructure (identity, payments, data-exchange, government cloud) that converts raw data into a production factor.

Why it matters. Data is the only AI input where latecomer nations can hold *absolute* advantages: population-scale, domain-specific, multilingual data that frontier labs cannot synthesize. India's DPI stack (Aadhaar, UPI at ~20 billion monthly transactions, DEPA consent architecture) demonstrates that public data rails can create adoption surfaces no private market would build.

Economic logic. Data is non-rival but excludable, exhibits steep quality-over-quantity returns, and is worthless without interoperability — hence infrastructure and data belong in one account: rails determine whether data assets have positive marginal product. National accounts book data at zero; GAIN capitalizes it using market-equivalent licensing values and shadow prices from data-acquisition deals.

Measurement methodology. A national data balance sheet: inventory of machine-readable, legally usable datasets by domain (health, finance, agriculture, language, geospatial, administrative), valued at replacement/licensing cost; DPI coverage rates (identity, payments, data-exchange penetration); interoperability score (share of government data exposed via standard APIs); and a data-governance usability index (can value be extracted lawfully — neither data anarchy nor data paralysis).

Core indicators. Capitalized national data assets (% GDP); DPI penetration (population share with digital ID, digital payments, consented data-sharing); open-data API coverage; sovereign language-corpus scale (tokens of high-quality national-language text); health/agriculture/finance dataset readiness scores; cross-border data-flow balance.

Data sources. National data-protection authorities; DPI operators (NPCI-type entities); World Bank GovTech and ID4D datasets; open-data portals; data-marketplace price observations.

Policy implications. High-DPI nations (India, Singapore, Korea) should weaponize the asset: mandate API-exposure of public data, build sovereign domain datasets where they hold global scale advantages. Low-DPI producers (the U.S.) face the opposite risk — world-class models with no public rails for population-scale service delivery.

3.5 Pillar P4 — Innovation Capital

Definition. The nation's capacity to create new AI artifacts and firms: research output that matters, models that are used, startups that scale, and intellectual property that earns.

Why it matters. Innovation capital determines a nation's *price-setting* position. Producers set prices for the world; consumers take them. The 23:1 U.S.–China private investment ratio and 10:1 ratio in newly funded AI companies measure exactly this divide.

Economic logic. AI innovation departs from patent-era economics: key artifacts are open-weighted (priced at zero, valuable as ecosystem gravity), the binding inputs are compute and talent rather than booked R&D, and value concentrates in a power-law tail of firms. Measurement must therefore weight *usage and capability* of artifacts over counts of disclosures — a cited preprint, a downloaded model, and a revenue-generating startup are the units, not patents.

Measurement methodology. Three sub-accounts: (a) Research capability — field-weighted citation impact, top-venue share, frontier-model capability gap (benchmark distance of best domestic model from world frontier); (b) Enterprise formation — AI startups founded/funded, graduation rates to scale ($100M+ revenue), exit value retained domestically; (c) Artifact influence — downloads/dependencies of domestically produced models and frameworks, weighted by usage.

Core indicators. Frontier capability gap (percentage points from world-best on composite benchmarks); newly funded AI companies per million population; AI venture capital (% GDP); domestic-model usage share in domestic API/token consumption; field-weighted AI citation impact; AI IP receipts (licensing income, % services exports).

Data sources. AI Index and Epoch capability registries; Crunchbase/PitchBook/Tracxn; Hugging Face and package-registry telemetry; patent and licensing-receipt data (used with the caveats of Part 1).

Policy implications. The capability-gap indicator disciplines sovereign-model programs: if the gap is widening despite spending, the binding constraint is upstream (P1/P2), and the rational policy is fast-follow fine-tuning on open weights, not frontier training runs.

3.6 Pillar P5 — Diffusion Capital

Definition. The nation's demonstrated capacity to absorb AI into actual production: firm adoption depth, government deployment, household use, and the organizational change that converts tools into output.

Why it matters. Diffusion, not invention, is where GPTs pay. The economic history of electricity and IT is unambiguous: nations that diffused fastest, not those that invented first, captured the productivity dividend. Stanford's Adoption Monitor shows even U.S. workplace adoption recently *reversing* — evidence that diffusion is fragile and cannot be assumed.

Economic logic. Diffusion is the Φ multiplier on the entire capacity stack: a nation operating at $\Phi = 0.3$ wastes 70% of its compute, talent, and data investment. It is governed by complementary-investment economics (the J-curve): organizational redesign is the gating factor, and SMEs — lacking redesign capacity — are systematically late, which is why the *depth distribution* of adoption (production-grade vs pilot) matters more than headline adoption rates.

Measurement methodology. A three-layer adoption census: firms (share using AI in production processes, weighted by employment and depth tier: experimenting / deployed / process-redesigned), government (share of high-volume public services with AI in the delivery path, with quality auditing), households (regular-use share, frequency-weighted). Supplemented by *intensity telemetry*: national API/token consumption per worker — the AI-era equivalent of electricity consumption per capita, and far harder to game than surveys.

Core indicators. Token consumption per worker; firm adoption depth index (production-deployment share, employment-weighted); SME adoption gap (large-firm minus SME adoption rates); government AI deployment share; automation/augmentation usage ratio at national level (importing the Anthropic Economic Index construct — because the Canaries evidence shows the *character* of usage predicts labor outcomes); AI-attributed process-cost reductions in benchmark sectors.

Data sources. Model-provider aggregate telemetry (negotiated national statistics partnerships — the single highest-value new data source available to statistical agencies); business surveys (Yotzov et al. model); government service audits; VAT/e-invoice data revealing AI-service purchases by firm size.

Policy implications. Diffusion is the highest-leverage, lowest-glamour pillar. The SME gap indicator tells governments where the multiplier is dying; token-per-worker tells them whether "adoption" is real usage or press releases.

3.7 Pillar P6 — Institutional Capital

Definition. The trust, governance quality, regulatory throughput, safety capability, and societal license that determine whether AI deployment is fast, durable, and politically sustainable.

Why it matters. Institutions set Ψ — the term that can zero out everything else. Public backlash, scandal-driven moratoria, or paralytic regulation can halt diffusion regardless of capacity; conversely, trusted certification regimes (Singapore's AI Verify) accelerate adoption by collapsing firms' legal uncertainty.

Economic logic. Governance is not a brake on the AI economy; it is a *throughput technology*. The correct measure is neither regulation quantity nor laissez-faire purity but regulatory velocity-with-trust: how fast can a high-stakes AI application get to deployment with public legitimacy intact? Institutional capital also includes state capacity to *evaluate* AI (sovereign testing and safety institutes) — a nation that cannot independently assess models must take vendors' word, which is a sovereignty deficit.

Measurement methodology. Composite of: public-trust surveys (trust in AI systems and institutions deploying them); regulatory throughput audit (median days from application to approval in regulated AI use cases — health, finance, transport); incident-and-redress audit (existence and speed of liability and recourse mechanisms); sovereign evaluation capability (does a national institute independently test frontier systems); workforce-transition coverage (share of displaced workers reached by reskilling/social insurance within 6 months).

Core indicators. AI public-trust index; regulatory throughput (median approval days, by sector); AI incident rate and redress latency; safety-institute capability tier; transition-coverage ratio; AI governance exports (other nations adopting your standards — a real, monetizable export).

Data sources. Edelman/Ipsos-style trust surveys nationalized; regulator administrative data; OECD.AI incident database; ILO transition-program data.

Policy implications. This account converts "responsible AI" from rhetoric to a managed KPI: a nation can see whether trust is rising in step with deployment, and act before the gap becomes backlash. Backlash is a *leading* indicator of diffusion collapse (Part 5).

3.8 Pillar P7 — Realized Returns

Definition. The economic value a nation actually harvests from its AI capacity: productivity gains, consumer surplus, fiscal revenue, export earnings, new-industry formation — and their distribution across people, firms, and regions.

Why it matters. Capacity without returns is a museum. The entire framework exists to maximize this account; measuring it separately is what makes the Conversion Ratio — returns relative to capacity — computable, and prevents nations from mistaking GPU purchases for economic strategy.

Economic logic. Returns arrive through four channels with different time constants and visibility: cost reduction in existing production (fast, partially visible in profits); consumer surplus (fast, invisible to GDP — requiring GDP-B-style measurement); new products/industries (slow, visible); and terms-of-trade gains from AI exports (slow, visible in the external account). Distribution belongs in the account because concentrated returns are politically self-limiting: they erode P6 and choke the flywheel (Part 7).

Measurement methodology. Four ledgers: (a) Productivity ledger — AI-attributed TFP and labor-productivity contributions estimated from firm-level adoption-performance microdata (difference-in-differences on adopters vs matched non-adopters), aggregated nationally; (b) Surplus ledger — GDP-B-style annual valuation of free/cheap AI services via choice experiments; (c) External ledger — the **AI Current Account**: exports minus imports of AI chips, compute services, model/API access, AI software, and AI-embodied services, plus the **AI Terms of Trade** index; (d) Distribution ledger — share of AI-attributed income gains accruing to labor; regional concentration (Gini of AI activity across subnational units); fiscal yield (AI-attributable tax revenue).

Core indicators. AI-attributed productivity contribution (pp/yr); consumer surplus (% GDP); AI Current Account balance (% GDP); AI Terms of Trade index; labor share of AI gains; regional AI-activity Gini; AI fiscal yield.

Data sources. Statistical-agency firm microdata linked to adoption telemetry; customs and balance-of-payments data with new AI product codes (a concrete proposal: an HS/EBOPS extension for AI services, analogous to the ICT goods classification); GDP-B survey instruments; tax administration data.

Policy implications. The Returns account is the minister of finance's page. If P1–P6 scores rise for three years while P7 stagnates, the nation is accumulating stranded capacity — and the framework points to which complement is missing.

3.9 What is genuinely new here

Five constructs in GAIN exist in no current framework: (1) the **capacity/returns separation** and the Conversion Ratio derived from it; (2) the **AI Current Account and AI Terms of Trade** — a full external account for AI value flows; (3) **token-consumption-per-worker** as a hard, telemetry-based diffusion metric replacing self-reported adoption; (4) the **apprenticeship-rung index**, nationalizing Stanford's canary insight into a monitored pipeline-integrity indicator; and (5) the **national data balance sheet**, capitalizing data as an asset class. Together they convert AI measurement from a research observatory into an accounting system a treasury can run.

Part 4. The National AI Economy Scorecard

4.1 Architecture

The scorecard operationalizes the seven NAIA accounts into eight scored dimensions (Realized Returns splits into its productivity and wealth components, which behave differently and deserve separate visibility). Each dimension is scored 0–100 from three to five quantitative indicators; dimensions aggregate into the headline **GAIN Score**, plus two diagnostic composites and the **Conversion Ratio**.

Dim	Dimension	NAIA account	Weight	Anchor indicators (top three)
D1	Compute Capital	P1	12%	AI compute stock per $bn GDP; compute self-sufficiency ratio; energy pipeline (GW)
D2	AI Human Capital	P2	16%	Talent density per 10k workers; net talent balance; apprenticeship-rung index
D3	AI Innovation Capital	P4	14%	Frontier capability gap; funded AI startups per million pop.; AI VC % of GDP
D4	AI Adoption Intensity	P5	16%	Token consumption per worker; firm production-deployment share; SME adoption gap (inverted)
D5	AI Productivity Impact	P7a	10%	AI-attributed TFP contribution; AI-attributed labor-productivity growth; sectoral cost reductions
D6	AI Public Infrastructure	P3	8%	DPI penetration; capitalized data assets % GDP; open-data API coverage
D7	AI Trust & Governance	P6	9%	Public-trust index; regulatory throughput; sovereign evaluation capability
D8	AI Wealth Creation	P7b–d	15%	AI Current Account % GDP; AI sector GVA % GDP; labor share of AI gains

The weighting logic is explicit and defensible before a peer-review panel: the two largest weights go to **Human Capital and Adoption Intensity (16% each)** because the historical record of general-purpose technologies shows absorption capacity, not invention, determines who captures the dividend; **Wealth Creation (15%)** ranks third because realized, retained value is the framework's purpose; **Innovation (14%)** and **Compute (12%)** follow as the principal capacity differentiators; **Productivity (10%)** is deliberately moderated because of J-curve measurement lag; **Trust (9%)** and **Public Infrastructure (8%)** are the multipliers — smaller in weight, but protected from irrelevance by the aggregation method below.

4.2 Formulas

Step 1 — Indicator normalization. Each indicator x is min–max normalized against a five-year rolling global panel, winsorized at the 5th/95th percentiles to prevent outlier compression:

> $I = 100 \times (x - x_5) / (x_{95} - x_5)$, clamped to [0, 100]

Directionally negative indicators (e.g., SME adoption gap, supplier-dependence HHI) are inverted ($100 - I$).

Step 2 — Dimension scores. Each dimension is the arithmetic mean of its normalized indicators (equal weights within dimension, revisited tri-annually by a methods board).

Step 3 — Headline aggregation (weighted geometric mean). Because the underlying accounts are economic complements, the GAIN Score uses a weighted geometric mean, which structurally penalizes imbalance:

> $\mathbf{GAIN = \Pi\, D_i{}^{\wedge}w_i = exp(\, \Sigma\, w_i \cdot ln\, D_i\,)}$, with $\Sigma w_i = 1$ and D_i floored at 5

A nation scoring 90 on compute and 10 on trust does not average to 50 — it scores 38. This single design choice encodes the framework's central economics: AI capacity factors multiply, they do not add. (An arithmetic-mean variant is published alongside for transparency; the gap between the two is itself reported as the **Imbalance Penalty**.)

Step 4 — Diagnostic composites.

- **Capacity Index (CI):** geometric composite of D1, D2, D3, D6, D7 (weights renormalized) — what the nation *has*.
- **Returns Index (RI):** geometric composite of D5, D8 (weights renormalized) — what the nation *gets*.
- **Conversion Ratio = RI / CI.** Below ~0.75: stranded capacity (assets without harvest). Near 1.0: balanced. Above ~1.15: over-extraction — harvesting imported capability without building any, the consumer trap.
- **GAIN Momentum:** three-year annualized change in GAIN, reported beside the level — because trajectory, not position, is the policy variable.

Step 5 — Maturity tiers. Frontier (GAIN ≥ 80) · Advanced (65–79) · Emerging (50–64) · Foundational (35–49) · Nascent (< 35).

4.3 Worked example: calculating India's GAIN Score

Illustrative 2026 dimension scores (calibrated to the Part 2 evidence; presented to demonstrate mechanics):

Dim	Dimension	Score	Weight	Basis (abridged)
D1	Compute Capital	28	12%	~38–58k public-program GPUs against a 100k target; very low stock per capita; strong subsidized-price access (₹65/hr); large energy pipeline but grid constraints

Dim	Dimension	Score	Weight	Basis (abridged)
D2	AI Human Capital	62	16%	Largest AI-skilled pool outside U.S./China; high T2/T3 volume, thin T1 frontier tier; persistent net emigration of top researchers partially reversing
D3	AI Innovation Capital	38	14%	Vibrant applied-startup scene; no frontier model; AI VC % GDP low; rising open-source contribution
D4	AI Adoption Intensity	44	16%	IT-services and fintech adoption strong; SME gap very large; token-per-worker far below OECD
D5	AI Productivity Impact	35	10%	Early; concentrated in exported services rather than domestic production
D6	AI Public Infrastructure	78	8%	World-leading DPI: Aadhaar ~1.4B, UPI ~20B monthly transactions, DEPA; data-asset capitalization unrealized
D7	AI Trust & Governance	52	9%	High public optimism toward AI; regulatory throughput untested; evaluation capability nascent
D8	AI Wealth Creation	41	15%	Substantial AI-embodied services exports; negative AI hardware current account; low domestic capture of model-layer value

Applying Step 3:

> ln-weighted sum = 0.12·ln28 + 0.16·ln62 + 0.14·ln38 + 0.16·ln44 + 0.10·ln35 + 0.08·ln78 + 0.09·ln52 + 0.15·ln41 = 3.792 **GAIN(India, 2026) = $e^{3.792}$ ≈ 44 — Foundational tier**

The arithmetic variant yields 46.2, so the **Imbalance Penalty is ≈ 2 points** — modest, indicating India's problem is *level*, not lopsidedness. The diagnostics: **Capacity Index ≈ 47, Returns Index ≈ 38, Conversion Ratio ≈ 0.82** — mild stranded capacity, concentrated in unmonetized DPI and talent. The single highest-leverage result: D1 (compute, 28) is India's binding factor; because aggregation is geometric, raising D1 from 28 to 45 alone lifts GAIN by roughly 2.6 points — more than any equivalent effort elsewhere. The scorecard does not merely rank India; it prices its bottleneck.

4.4 Illustrative 2026 league table

Country	GAIN	Capacity	Returns	Conversion	Tier and one-line diagnosis
United States	88	91	86	0.95	Frontier — full stack; watch trust and distribution
China	80	82	77	0.94	Frontier — diffusion superpower; frontier-compute constrained

Country	GAIN	Capacity	Returns	Conversion	Tier and one-line diagnosis
Singapore	73	75	70	0.93	Advanced — orchestration model at physical limits
South Korea	69	72	64	0.89	Advanced — component power converting to sovereign capability
United Kingdom	67	74	58	0.78	Advanced — capacity rich, capture poor: the leakage case
France	62	66	55	0.83	Emerging→Advanced — Europe's producer candidate
UAE	60	63	55	0.87	Emerging — platform rents real; ecosystem conversion unproven
Canada	58	67	47	0.70	Emerging — stranded capacity flagship: invents, others harvest
Saudi Arabia	48	52	42	0.81	Foundational — capital-led buildout ahead of talent base
India	44	47	38	0.82	Foundational — compute-bound diffusion giant (Part 6)

Two readings of this table matter more than the ranking itself. First, the **Conversion column re-orders the world**: the UK and Canada — celebrated in conventional indices — appear here as the leading cases of capacity leakage, while the UAE's modest GAIN level conceals an efficient conversion machine. Second, **momentum diverges from level**: on three-year trajectory, India, Saudi Arabia, and the UAE are the fastest risers; the UK is flat; Canada is the only profiled nation whose Conversion Ratio is *falling*. A government using this scorecard annually is therefore managing three numbers — level, conversion, momentum — not one.

Part 5. The AI Value Creation Model: How AI Becomes National Wealth

5.1 The value chain, stage by stage

National AI wealth is created along an eight-stage chain. Each stage has a distinct economic function, a characteristic time constant, a leakage mode (how value escapes the nation), and a government intervention with measurable leverage.

> AI Research → AI Infrastructure → AI Startups → AI Adoption → Productivity → New Industries → Export Capacity → National Wealth

Stage	Economic function	Time constant	Dominant leakage mode	Highest-leverage intervention
1. Research	Creates capability options (real options on future industries)	5–15 yrs	Talent emigration; foreign acquisition of labs	Fund people not projects; compute-for-researchers; repatriation packages
2. Infrastructure	Converts energy + capital into usable compute and rails	2–5 yrs	Foreign ownership of domestic capacity; rented sovereignty	Power-grid throughput; permitting velocity; ownership-tier rules; subsidized access pricing
3. Startups	Converts capability into firms and priced products	3–7 yrs	Flip-and-exit to foreign acquirers; scale-up capital flight	Late-stage domestic capital pools; procurement-as-first-customer; talent visas
4. Adoption	Diffuses capability into the existing economy	3–10 yrs	Pilots that never reach production; SME exclusion	SME diffusion programs; government-as-exemplar; standards that cut integration cost
5. Productivity	Converts usage into output per worker	5–12 yrs (J-curve)	Gains absorbed by complementary-investment costs; mismeasured and politically unclaimed	Management-capability programs; process-redesign tax treatment (expense → capital)
6. New industries	Creates products impossible before AI	7–15 yrs	First markets captured by foreign platforms	Regulatory sandboxes; sovereign demand for AI-native services (health, agriculture, education)
7. Export capacity	Sells AI inputs/outputs abroad; sets prices instead of taking them	5–15 yrs	Terms-of-trade deterioration; brand absence	Trade missions for AI services; AI product codes in trade policy; governance exports
8. National wealth	Accumulates returns as fiscal yield, wages, and	10–25 yrs	Offshore profit shifting;	Distribution monitoring; sovereign co-investment

Stage	Economic function	Time constant	Dominant leakage mode	Highest-leverage intervention
	reinvestable capital		concentration that erodes legitimacy	vehicles; reinvestment mandates

Three structural properties of the chain determine strategy. First, it is **leaky at every joint** — the UK loses between stages 1 and 3 (science to scale-up), Canada between 1 and 2 (research to infrastructure), India between 4 and 5 (adoption to measured productivity), the Gulf states risk losing between 2 and 3 (infrastructure to indigenous firms). National strategy is therefore less about pushing harder at any stage than about **sealing one's characteristic leak**. Second, the chain is **not strictly sequential**: adoption (4) can run on imported capability while research (1) matures — this is precisely India's play — but skipping stages permanently caps the Conversion Ratio, because price-setting power lives upstream. Third, **time constants lengthen downstream**: a government can buy stage 2 in three years, but stages 5–8 are decade-scale, which is why electoral cycles systematically over-invest in visible infrastructure and under-invest in diffusion and distribution.

5.2 Where governments should intervene: the leverage map

Not all public dollars are equal along the chain. Synthesizing diffusion economics and the historical GPT record, the model ranks intervention leverage by *marginal national return per public dollar*, conditional on a nation's archetype:

1. **For Consumers and Ecosystem Builders (India, Canada, most nations):** highest leverage at stages 4 and 2 — diffusion programs and compute access. The productivity dividend of adopting world-frontier capability already exists and is unclaimed; the binding constraints are integration cost and compute price. Frontier research subsidies, by contrast, buy lottery tickets against incumbents holding hundred-fold capital advantages.
2. **For Platforms (UAE, Saudi Arabia):** highest leverage at stage 3 — converting hosted infrastructure into indigenous firms before the compute rents commoditize. Compute is becoming abundant; ecosystems remain scarce.
3. **For Producers (U.S., China):** highest leverage at stages 5 and 8 — the frontier is already private; the public problems are grid throughput, diffusion into the long tail, and distribution legitimacy.

The general rule the model yields: **public money should buy whatever the private market structurally underprices in that nation** — basic research options, diffusion externalities, SME integration, trust infrastructure — and should never crowd into whatever global private capital is already oversupplying (frontier training runs, in 2026, being the canonical example).

5.3 Leading and lagging indicators

The chain implies a strict temporal ordering of signals. Governments managing only lagging indicators are steering by the rear-view mirror; the NAIA's operational value is that every account contains forward-looking series.

Horizon	Leading indicators (move 2–5 yrs ahead)	Lagging indicators (confirm, 5–10 yrs behind)
Capability	Net talent balance; compute pipeline (GW permitted); research citation momentum; frontier capability gap trend	Patent grants; university rankings; Nobel-class recognition
Commercial	New AI-firm formation rate; AI VC commitments; token consumption per worker growth; SME adoption-gap closure rate	AI sector GVA; employment in AI occupations; unicorn counts
Macro	Apprenticeship-rung index; automation/augmentation usage ratio; AI input price trends (GPU-hour, token prices); regulatory throughput times	TFP growth; labor share; median wages; GDP composition shifts
Political economy	Public-trust index trend; incident-and-redress latency; regional concentration of AI activity	Backlash legislation; moratoria; electoral realignment around AI

Two leading indicators deserve singling out as the system's smoke detectors. The **apprenticeship-rung index** (early-career hiring in exposed occupations — Stanford's canary, nationalized) leads both the skills crisis and the political crisis by half a decade: a nation whose young cannot enter cognitive careers will register the damage in trust surveys and ballot boxes long before it appears in unemployment statistics. And the **automation/augmentation ratio** leads the *character* of the transformation: the Canaries evidence already shows automation-skewed usage predicts employment decline where augmentation-skewed usage does not — making this ratio the earliest observable signal of whether a nation is walking into the Turing Trap or around it.

Part 6. India: A Full GAIN Case Study

6.1 Where India stands today

Run through the GAIN lens, India in 2026 is the world's most asymmetric AI economy: **first-quintile human capital and digital public infrastructure mounted on last-quintile compute**, with everything else in between. Its GAIN Score of ≈44 (Foundational tier, Part 4.3) understates its strategic position, because GAIN level is a stock measure and India's distinguishing features are flows: the fastest-compounding adoption surfaces on earth, and a state that has — unusually — already demonstrated it can build population-scale digital systems that work.

The honest one-line verdict the framework produces: **India today is the world's largest AI Consumer, running the most credible Consumer→Ecosystem-Builder transition anywhere, gated almost entirely by compute capital and diffusion depth rather than by talent or rails.**

Strengths: three assets of genuinely global rank

The talent reservoir (D2 = 62). India holds the largest AI-capable workforce outside the U.S. and China: millions of software developers, the largest national cohort of AI-project contributors on GitHub, and an engineering-education system producing at a scale no other democracy matches. Critically, this talent is *services-deployed* — already embedded in the global enterprise workflows (through TCS, Infosys, and the global capability centers of nearly every Fortune 500 firm) into which AI is now being integrated. India does not need to win frontier research to monetize this; it sits at the integration layer of world AI adoption, the layer where most of the realized value of the next decade will be created.

The DPI rails (D6 = 78). Aadhaar (~1.4 billion identities), UPI (on the order of 19–20 billion transactions monthly, the world's largest real-time payment fabric), DigiLocker, ONDC, and the DEPA consent architecture constitute the world's most advanced population-scale digital public infrastructure. In GAIN terms this is pre-built *diffusion surface*: every AI service that needs identity, payment, and consented data access can reach a billion users through rails that already exist. No other nation — including the U.S. — possesses this. It is also an unrealized *data balance sheet*: multilingual, domain-diverse, population-scale data assets that remain largely uncapitalized.

The export beachhead (D8 component). India already runs a large surplus in AI-*embodied* services — the export of human-plus-AI cognitive work through its IT sector. As enterprise AI shifts from chat tools to agentic process transformation, India's services giants are positioned to sell the transformation itself. The AI Current Account decomposition matters here: India is deeply negative on AI *hardware* and model-layer *royalties*, strongly positive on AI-embodied *services* — a terms-of-trade profile resembling a refiner that imports crude.

Weaknesses and structural bottlenecks

Compute poverty (D1 = 28) — the binding constraint. The IndiaAI Mission has taken common compute from zero to roughly 38,000–58,000 GPUs (with a 100,000 target by end-2026) at subsidized prices of ₹65–150/hour — genuinely effective policy. But the absolute scale remains one to two orders of magnitude below requirement for a nation of India's ambition: single U.S. hyperscaler campuses now exceed India's entire national stock. Per-capita compute is among the lowest of any major AI economy. Behind silicon sits power: AI-grade data centers need firm, cheap, growing electricity, and India's grid — though expanding renewables aggressively — faces transmission, reliability, and industrial-tariff constraints that gate the buildout more than capital does.

The frontier gap (D3 = 38). No Indian frontier model; modest AI VC (% of GDP among the lowest in the profiled ten); a startup ecosystem vibrant at the application layer but thin at the infrastructure and model layers; research output large in volume, weak in field-weighted impact. The T1 (frontier researcher) tier of the talent pyramid is disproportionately employed abroad.

The diffusion gap beyond the formal sector (D4 = 44). Indian AI adoption is deep in IT services, fintech, and the digital-native economy — and shallow almost everywhere else. The MSME sector (~110+ million workers) and agriculture (~40%+ of employment) are barely touched. Token-consumption-per-worker, the framework's hard diffusion metric, would place India far below OECD levels. This is the gap between *having* rails and *running trains* on them.

State-capacity variance (D7 = 52). India's policy direction is sound — the IndiaAI Mission's seven pillars map remarkably well onto the NAIA accounts — but execution oscillates: procurement delays, regulatory uncertainty (the evolving DPDP data regime), and the chronic announcement–implementation gap. Trust in AI among the public is among the world's highest (a real asset), but institutional evaluation capability and redress mechanisms are nascent.

6.2 Scenario analysis: India 2030 and 2035

The scenarios below vary four policy variables the framework identifies as decisive for India specifically: compute buildout (D1), MSME/agriculture diffusion (D4), frontier-adjacent capability (D3), and talent retention (D2). All GAIN projections are illustrative model outputs holding other nations' trajectories at consensus paths — meaning *relative* position shifts are conservative, since leaders are also moving.

Scenario 1 — Conservative: "Services absorber" (probability-weighted baseline if current execution patterns persist)

Compute reaches the 100k-GPU public target late and plateaus; private data-center investment grows but foreign-owned; diffusion stays formal-sector; no sovereign model of consequence; talent emigration at the top tier continues. India rides global AI as the world's integration workshop — substantial, real, and structurally subordinate: AI terms of trade stay refiner-like, price-taking on models and hardware.

> **GAIN 2030 ≈ 50 (Emerging, lower bound). GAIN 2035 ≈ 55.** Conversion Ratio stagnant ~0.8. India in 2035: the indispensable back office of the AI economy — prosperous at the integration layer, absent from the price-setting layers, with an MSME economy that largely missed the dividend.

Scenario 2 — Accelerated: "DPI-powered diffusion state" (the plausible high path with sustained execution)

Compute scales to several hundred thousand publicly accessible accelerators by 2030 via public-private builds anchored on dedicated renewable/nuclear power corridors; DPI rails are deliberately weaponized for diffusion — AI-on-UPI-scale deployments in credit, agriculture advisory, diagnostics, and education reach hundreds of millions; an "ONDC for AI" standardizes MSME integration, halving the SME adoption gap; sovereign Indic-language models reach near-frontier capability in their domains via fast-follow on open weights; GCC deepening turns talent emigration into circulation. Productivity effects become measurable economy-wide in the early 2030s.

> **GAIN 2030 ≈ 57 (Emerging, upper half). GAIN 2035 ≈ 66 (Advanced).** Conversion Ratio rises toward 0.95 as DPI and talent monetize. India in 2035: the world's third AI economy by realized returns, the global reference case for diffusion-led (rather than frontier-led) AI development, and the standards-setter for the Global South's AI adoption — governance and DPI exports becoming a real current-account line.

Scenario 3 — AI Superpower: "Full-stack democracy" (possible; requires three simultaneous breaks from precedent)

Everything in Scenario 2, plus: (a) an energy break — dedicated AI power buildout on the scale of tens of GW, making India a cost-competitive compute *exporter* to Asia and Africa; (b) a capital break — domestic late-stage capital pools (sovereign co-investment, pension reform) that keep Indian AI firms Indian through scale-up; (c) a frontier break — at least one Indian lab at the global frontier in a consequential domain (most plausibly: multilingual/agentic systems for population-scale services, where India's data advantage is absolute rather than relative). The apprenticeship-rung is actively defended, making India the demographic counterweight to an aging, automation-anxious West.

> **GAIN 2030 ≈ 62 (borderline Advanced). GAIN 2035 ≈ 74 (Advanced, approaching Frontier).** India in 2035: the third pole of the AI world — not by matching U.S./China frontier compute, but by owning the diffusion-and-scale layer of the global AI economy and pricing it. AI Current Account positive across services *and* compute hosting; terms of trade transformed.

What separates the scenarios

The framework reduces the difference between Scenario 1 and Scenario 3 to four measurable variables, in causal order of urgency: **megawatts** (firm power dedicated to compute — the true unit of Indian AI policy, more binding than GPUs); **the SME adoption gap** (the entire difference between a services enclave and an economy-wide dividend); **late-stage capital domesticity** (who owns Indian AI firms at maturity); and **the T1 talent balance** (whether the top thousand minds circulate or simply leave).

Every one of these is observable quarterly in the NAIA. None requires waiting for GDP statistics to know whether the strategy is working — which is precisely the point of the framework.

Part 7. The AI Economy Flywheel

7.1 The compounding loop

National AI advantage is not a position but a *rotation*. The GAIN Flywheel describes the eight-station loop through which AI capacity compounds — each station converting one form of capital into the next, with national wealth feeding back into the loop's first station:

> **Talent → Compute → Startups → Adoption → Productivity → Wealth → Reinvestment → Talent** (and around again)

The conversions, station by station: *Talent → Compute* — skilled people justify and absorb infrastructure investment (compute without users is stranded capital; this is why talent precedes compute in the loop, a sequencing most national strategies invert). *Compute → Startups* — cheap, accessible compute collapses the cost of founding AI-native firms (the IndiaAI ₹65/hour subsidy and Canada's Compute Access Fund are exactly this conversion, engineered). *Startups → Adoption* — domestic firms localize global capability into products incumbents and SMEs can actually buy, in local language, regulation, and workflow. *Adoption → Productivity* — usage plus organizational redesign yields output per worker, after the J-curve's complementary-investment delay. *Productivity → Wealth* — productivity becomes profits, wages, and fiscal revenue; consumer surplus accrues invisibly but politically matters. *Wealth → Reinvestment* — the loop's governance moment: gains are recycled into the cycle's inputs, or they are not. *Reinvestment → Talent* — funded universities, research careers worth staying for, and an economy where AI skills command returns — closing the loop and accelerating it.

The flywheel's defining property is that **rotation speed compounds**: each full cycle leaves every station larger than the last, which is why small early leads (the U.S. post-2012, deep-learning era) become enormous and why late starters cannot catch up by matching current investment — they must match *cumulative rotations*, or find a shortcut station (India's DPI is precisely such a shortcut: pre-built adoption surface that lets the wheel skip years of diffusion friction).

7.2 Where nations fail: the eight stall points

Each station has a characteristic failure mode — and each failure has a name in the 2026 landscape:

Stall point	Failure mode	Exhibit
Talent → Compute	Train talent, deny it compute; researchers emigrate to where the GPUs are	Europe pre-2024; India pre-IndiaAI Mission
Compute → Startups	Build sovereign compute used at 30% by incumbents and ministries; no founder access	Risk now facing several sovereign-compute programs, incl. Gulf builds

Stall point	Failure mode	Exhibit
Startups → Adoption	Startups form, but domestic enterprises buy foreign platforms; local firms flip abroad	UK's chronic scale-up exodus; Canada's acquisition drain
Adoption → Productivity	Pilots without process redesign — the J-curve becomes a flat line; "AI theater"	The global enterprise default, 2023–26; visible in adoption-reversal survey data
Productivity → Wealth	Gains real but captured offshore or by a narrow capital stratum; nation hosts the work, books little of the income	Profit-shifting in digital services; the platform-economy precedent
Wealth → Reinvestment	Gains consumed (tax cuts, transfers, rents) rather than recycled into talent and infrastructure	The resource-curse pattern; the open risk for compute petrostates
Reinvestment → Talent	Money recycled, but into buildings and programs rather than people and careers; brain drain persists	The lab-without-salaries failure common to mid-tier science systems
Cross-cutting	Trust collapse: an incident-and-backlash spiral seizes the whole wheel regardless of station strength	The scenario every Part 8 trust instrument exists to prevent

The diagnostic discipline the flywheel imposes: **a nation should locate its slowest station before spending anywhere else**, because flywheel output is gated by the weakest conversion, not the strongest station. This is the same complementarity economics as the geometric mean in Part 4 — the scorecard and the flywheel are the static and dynamic views of one model.

7.3 How successful nations sustain rotation

Three sustained-rotation patterns emerge from the profiled ten. The **U.S. model — private flywheel, public bearings**: the loop runs on private capital and labor markets; the state's role is keeping bearings frictionless (research funding at the talent station, permitting and grid at compute, antitrust and immigration as lubricants — the last now visibly degrading). The **China model — state-forced rotation**: where a station underperforms, the state forces conversion (guidance funds at startups, AI+ mandates at adoption, procurement at productivity); costlier per rotation, but immune to private-capital flight, and uniquely strong at the adoption station. The **Singapore model — precision flywheel**: a small wheel spun at maximum efficiency — every station instrumented, no station oversized, surplus exported as governance IP. The common invariant across all three: **the reinvestment station is governed deliberately** — by venture-market depth in the U.S., by state planning in China, by sovereign-fund discipline in Singapore. Nations fail at the wheel not mostly from poverty but from letting the Wealth → Reinvestment conversion happen by accident.

Part 8. The Government Playbook

8.1 Sequenced actions by pillar

The playbook below is archetype-neutral (Part 5.2 sets priorities by archetype); horizons reflect realistic institutional lead times. The unifying rule: **years 0–2 buy data and access, years 3–5 buy conversion, years 5–10 buy position.**

Pillar	Immediate (0–2 yrs)	Medium (3–5 yrs)	Long (5–10 yrs)
P1 Compute	Stand up the compute census; designate AI power corridors and fast-track permitting; negotiate sovereign cloud-access contracts; subsidize researcher/startup access pricing	Build or anchor 2–4 national-scale facilities with mandated open-access tiers; diversify accelerator supply; lock long-term power purchase agreements	Energy-position strategy (nuclear/renewable baseload for compute); regional compute-export capacity; participate in setting international compute-governance regimes
P2 Talent	Count the stock (talent census); frontier-talent retention packages (compute + salary + autonomy); fast AI visas; emergency AI-fluency programs for civil service	Rebuild the apprenticeship rung: subsidized early-career AI-adjacent placements; reform curricula from coding to AI-orchestration skills; diaspora circulation programs	Research careers competitive with industry; K-12 AI-native education; become a net talent importer in at least one AI specialty
P3 Data/DPI	Inventory and legally clear high-value public datasets; expose top-20 government data APIs; adopt consent architecture	Build sovereign domain corpora where the nation holds scale advantages (language, health, agriculture); capitalize data assets onto a published national balance sheet	Full DPI stack if absent (identity, payments, data exchange); cross-border data-flow treaties; data-asset licensing as fiscal revenue
P4 Innovation	Procurement-as-first-customer for domestic AI firms; regulatory sandboxes in two regulated sectors; fund fast-follow fine-tuning, not frontier vanity runs	Late-stage domestic capital pools (sovereign co-investment, pension allocation reform); university spin-out reform (IP terms, founder leave)	One credible frontier-adjacent lab in a domain of absolute national advantage; domestic acquisition capacity so exits recycle locally
P5 Diffusion	Government-as-exemplar: AI in three high-volume public services with published metrics; SME diffusion vouchers; token-consumption telemetry partnership with providers	National process-redesign program (management capability, not just tools); halve the SME adoption gap; sectoral diffusion compacts (agriculture, health, logistics)	AI-native public services as the default; diffusion institutions with the permanence of agricultural extension services
P6 Trust	Publish incident-and-redress framework; stand up sovereign evaluation capability; baseline public-	Regulatory throughput targets (median approval days) in health/finance/transport;	Export the governance stack (certification regimes, audit standards); embed AI legitimacy in

Pillar	Immediate (0–2 yrs)	Medium (3–5 yrs)	Long (5–10 yrs)
	trust survey	transition-coverage guarantee for displaced workers	constitutional-grade institutions
P7 Returns	Begin the four ledgers (productivity, surplus, external, distribution); add AI product codes to trade statistics	Publish the AI Current Account annually; distribution monitoring with regional dashboards; fiscal instruments tuned to AI value chains	Sovereign wealth participation in AI returns; rebalance taxation from labor toward AI-era bases as the labor share shifts

8.2 The national measurement cadence

What governments should measure monthly — the dashboard of fast-moving series: token consumption per worker (national aggregate and by sector); GPU-hour reference price and compute utilization in public facilities; the apprenticeship-rung index (early-career hiring in exposed occupations, the Canaries construct); AI job-posting share and its junior/senior split; the automation/augmentation usage ratio; new AI-firm registrations; public-service AI deployment metrics; and the incident log with redress latency.

What central banks should monitor: AI-attributed productivity contribution and its sectoral diffusion (for potential-output and r* estimation — an AI productivity acceleration changes the neutral rate); labor-share trajectory and wage dispersion between AI-leveraged and AI-displaced occupations (for the employment mandate and inflation dynamics); AI capital expenditure as a share of gross fixed capital formation and its financing structure (credit concentration in data-center lending is a nascent financial-stability exposure); AI-related equity concentration and collateral channels; and the AI import bill's pass-through to the current account and exchange-rate pressure — for compute-importing nations, GPU and API dependence is a terms-of-trade exposure exactly as oil was.

What ministries should monitor: Finance — the Conversion Ratio, AI fiscal yield, and the four Returns ledgers. Industry/Economy — SME adoption gap, sectoral deployment depth, startup formation and exit domesticity. Education/Labor — talent stock-and-flow table, apprenticeship-rung index, reskilling completion-to-employment rates. Energy — AI power-corridor pipeline (GW permitted/connected), compute facilities' grid queue position. Trade — AI Current Account, AI Terms of Trade, services-export AI intensity. Health/Agriculture/Interior — domain deployment metrics with quality audits. Digital/IT — DPI penetration, API coverage, sovereign corpus growth, evaluation-institute throughput.

What investors should monitor (the framework doubles as a country-allocation tool): GAIN Momentum over GAIN level — trajectory prices better than position; the Conversion Ratio — rising conversion signals an economy about to monetize visible capacity (the 2026 long candidates on this screen: India, UAE; the short: capacity-rich, conversion-flat incumbents); the compute-price-to-talent-density ratio as a venture-cost arbitrage signal; regulatory throughput times as the best single proxy for deployment-market openness; and the apprenticeship-rung and trust indices as early warning of the political risk that re-prices everything else.

8.3 Institutionalizing the system

The playbook's final recommendation is institutional: every nation serious about this agenda should charter a **National AI Accounts Unit** inside its statistical agency (not its innovation ministry — measurement must outlive strategy), publishing on a fixed calendar: monthly dashboard (the series above), quarterly NAIA accounts, and an annual GAIN report with international benchmarking. Three design rules keep it honest: telemetry over surveys wherever possible; methodology published to academic standard and governed by an external methods board on tri-annual revision cycles (the Stanford Lab, the OECD, and the IMF are the natural custodians of cross-country comparability); and a statutory firewall against the agency that runs AI strategy grading its own homework.

Part 9. Executive Summary

9.1 The question, answered

> **"What should governments measure if they want to build an AI economy rather than merely adopt AI?"**

Adopting AI requires measuring usage. Building an AI economy requires measuring **ownership of capacity, efficiency of conversion, and direction of flows** — seven accounts, three numbers, and one discipline:

Measure seven capitals, not one activity. Compute (the new energy), cognitive capital (the new labor), data and digital rails (the new land), innovation capability (the new entrepreneurship), diffusion depth (the new electrification rate), institutional trust (the new rule of law), and realized returns (the only reason any of it matters). These are the National AI Accounts (NAIA). A nation that tracks adoption alone is measuring its consumption of someone else's economy.

Manage three numbers, not a ranking. The **GAIN Score** (level — a weighted geometric composite, because AI capacities multiply rather than add); **GAIN Momentum** (trajectory — the only variable policy moves on electoral timescales); and the **Conversion Ratio** (returns relative to capacity — the difference between owning AI assets and monetizing them; the number that exposes the UK and Canada as leakage economies and the UAE as an efficient converter, whatever conventional rankings say).

Watch the border and the rung. The **AI Current Account** — net trade in chips, compute, models, tokens, and AI-embodied services — is the strategic balance of the 2030s; a nation that does not construct it is running industrial policy blind. And the **apprenticeship-rung index** — early-career hiring in AI-exposed occupations, generalizing Stanford's canaries finding of −3.8% annual contraction — is the earliest warning the system gives before a skills pipeline and its politics fail together, and the empirical tripwire for the Turing Trap: economies rewarding automation over augmentation register it here first.

9.2 The findings in brief

1. **First-generation AI measurement is an observatory, not an operating system.** Stanford's AI Economic Indicators — the field's best instruments — track effects (employment, adoption, macro ratios) in one country, with no accounts of capacity, ownership, trade, or distribution, and no embedded theory of action (Part 1).
2. **Nations are playing four different games.** Producers (U.S., China) sell the means of AI production; Platforms (UAE, Saudi Arabia) convert energy and capital into hosting rents; Ecosystem Builders (Singapore, Korea, Canada, France) compete on orchestration; Consumers — most of the world — take prices.

Strategy begins with knowing your game; the PPCE typology and two-axis model make position and trajectory measurable (Part 2).

3. **Complementarity is the central economics.** AI capacity factors multiply: compute without talent, adoption without redesign, capability without trust each yield approximately nothing. The framework encodes this in a geometric production function, a geometric scorecard, and a flywheel gated by its weakest conversion (Parts 3, 4, 7).
4. **Diffusion, not invention, pays the dividend** — the verdict of every general-purpose technology in history, weighted accordingly (16% adoption, 16% human capital) and measured with telemetry (token consumption per worker), not self-reported surveys (Parts 3–4).
5. **India is the decisive experiment.** GAIN ≈ 44 today — world-class talent and rails on last-quintile compute. Its three scenarios (≈55 / ≈66 / ≈74 by 2035) span the difference between the world's back office and the world's third AI pole, and the gap reduces to four measurable variables: megawatts, the SME gap, capital domesticity, and frontier-talent circulation (Part 6).
6. **The reinvestment station is where nations are won and lost.** Across every model that sustains the flywheel — American private markets, Chinese state forcing, Singaporean precision — the Wealth → Reinvestment conversion is governed deliberately. Nations rarely fail from poverty; they fail from letting compounding happen by accident (Part 7).

9.3 The GAIN Framework as scalable intellectual property

The framework is engineered for five lives beyond this report. As a **national competitiveness index**: an annual GAIN league table with the Conversion Ratio as its signature analytic — the number no existing index publishes. As a **policy framework**: the NAIA gives treasuries and statistical agencies an implementable accounting standard (chartered units, fixed cadences, methods-board governance), with the AI Current Account as its first concrete statistical innovation — a natural candidate for OECD/IMF working-group adoption, as the ICT classifications were a generation ago. As a **consulting methodology**: the diagnostic sequence — classify the archetype, score the eight dimensions, locate the binding constraint, find the flywheel stall, sequence the playbook — is a repeatable engagement architecture for governments and sovereign funds. As a **research institute's flagship model**: every construct is falsifiable and improvable — the weights, the exposure measures, the conversion elasticities between pillars are a decade-long empirical agenda, and the framework is designed to be wrong in public and revised on schedule, which is what made the System of National Accounts durable. And as an **annual report**: *The State of the AI Nations*, built on the accounts, anchored each year by the three numbers and the league table.

The closing argument is the opening one. In 1937, nations learned to measure the industrial economy only after it had collapsed beneath them. The AI economy is being assembled now, in plain sight, at documented speed — its capital accounts visible in grid queues and customs data, its labor effects already legible in the hiring of 22-year-olds, its trade flows already pricing in GPU-hours and tokens. The nations that prosper in it will be those that stopped measuring the economy they used to have and started keeping accounts of the one being built. **What gets measured gets governed; what gets governed gets compounded.** The instruments now exist. The only remaining question is which governments will be first to run them.

Appendix: Sources and Foundations

The framework synthesizes and extends: Stanford Digital Economy Lab, *AI Economic Indicators: June 2026 Update* (Research Note No. 1) — the Canaries Dashboard (with ADP Research), Takeoff Tracker, and Adoption Monitor, including Brynjolfsson, Chandar & Chen (2025), Eloundou et al. (2024), and the Anthropic Economic Index automation/augmentation constructs; Stanford HAI, *AI Index Report 2026* (investment, capability-gap, talent-migration, and consumer-surplus estimates); Brynjolfsson's research program on the productivity J-curve, GDP-B, and the Turing Trap; OECD digital-economy measurement frameworks and the OECD.AI policy observatory; World Bank digital development and ID4D/GovTech indicators; IMF productivity and labor analyses; WEF *Future of Jobs* and AI-sovereignty reports; the EU Digital Economy and Society Index (DESI) methodology; and national strategy documents and program data for the ten profiled countries, including the IndiaAI Mission, Canada's Sovereign AI Compute Strategy, Korea's AI Action Plan (2026), the UAE's Stargate program, Saudi Arabia's HUMAIN, and the UK and French national AI programs, current as of June 2026. Country scores and projections are the authors' inference-based calibrations for methodological illustration.