

The Abundance Dividend

What safe, broadly deployed AI is worth to humanity — and how to actually collect it

PREPARED FOR ECONOMISTS • STRATEGISTS • POLICYMAKERS • INVESTORS



FAIR Labs

FAIR ARTIFICIAL INTELLIGENCE RESEARCH LABS

July 2026
FL-2026-02

The Abundance Dividend: What safe, broadly deployed AI is worth to humanity — and how to actually collect it

The second volume of the Frontier Series prices the upside. Safe, broadly deployed AI could compress the cost of cognitive work the way electrification compressed the cost of power — but abundance is a dividend, not a default. It must be claimed through deployment, diffusion, and trust, and it is paid in proportion to all three. The report maps five abundance frontiers, diagnoses the diffusion bottleneck, argues that assurance is an economic input, and closes with a seven-point compact for collecting the dividend broadly.

REPORT CODE	FL-2026-02
PUBLICATION DATE	July 2026
PRIMARY AUDIENCES	Economists · Strategists · Policymakers · Investors
SERIES	FAIR Labs Frontier Series, Fair Artificial Intelligence Research Labs
SUGGESTED CITATION	<i>The Abundance Dividend: What safe, broadly deployed AI is worth to humanity — and how to actually collect it.</i> FAIR Labs Frontier Series, Vol. 02. Fair Artificial Intelligence Research Labs, 2026.
KEYWORDS	AI economics, abundance, productivity, diffusion, distribution, growth, public benefit

This report presents analysis and frameworks developed by FAIR Labs. Quantitative exhibits marked as illustrative or as FAIR Labs assessments are analytical constructs intended to support reasoning, not measurements. This report is for informational purposes and does not constitute legal, financial, or investment advice.

© 2026 Fair Artificial Intelligence Research Labs. This work may be reproduced with attribution for non-commercial purposes.

IN BRIEF • EXECUTIVE SUMMARY

The abundance is real. The dividend is optional.

Advanced AI keeps being described as a tide that lifts all boats. This volume drops the metaphor and reads the ledger. Cheap, capable cognition can underwrite the largest broad-based improvement in living standards since electrification — but the gains arrive only where systems are deployed, diffused beyond the sophisticated few, and trusted with work that matters. Abundance, in other words, is a dividend: real, large, and payable strictly in proportion to what societies build in order to collect it.

01 Abundance is a dividend, not a default.

Nothing about capable models guarantees broad benefit. The dividend is declared in the lab and collected in the economy — through deployment, diffusion, and trust. Each is a choice, and each is currently being fumbled somewhere.

02 The frontiers are concrete, and they are five.

Science, medicine, education, energy and materials, and state capacity hold most of the latent dividend. In four of the five, the binding constraint is already trust and integration, not model capability.

03 Diffusion, not invention, is the bottleneck.

Value scales with adoption depth: how many organizations, how deep into the workflow, how high the stakes entrusted. Frontier quality improves monthly; median adoption remains shallow, low-stakes, and concentrated among the already-sophisticated.

04 Trust is an economic input.

High-value domains admit AI by permission — from regulators, insurers, procurers, and publics. Assurance is the permission slip. Volume 01's capability-assurance gap is therefore not only a safety problem; it is unrealized growth.

05 Who collects is a design choice.

Left to defaults, the dividend concentrates — compute capital, data loops, and distribution moats see to that. Levers exist that widen collection without slowing the engine; the mistake is confusing distribution policy with capability policy.

06 Abundance has failure modes of its own.

Dependency without resilience, hollowed institutions, rebound congestion, cost-disease inversion, winner-take-most geography. None is a reason to slow down; all are reasons to steer. The closing compact assigns the steering.

CONTENTS

Inside this report

01	Cognition becomes a utility	5
02	Five abundance frontiers	7
03	The diffusion problem	10
04	The trust dividend	12
05	The distribution question	14
06	The failure modes of plenty	17
07	The Abundance Compact	19

01


Cognition becomes a utility

Every general-purpose technology has a moment when the thing it produces stops being an event and becomes an assumption. Light had it; calculation had it. Competent cognitive work is having it now — and economies organized around cognitive scarcity are about to discover what they had been paying for it.

The instructive precedents are not the machines but the prices. For most of recorded history, artificial light was a luxury metered in labor — hours of work for an evening of reading. Once cheap and ambient, it stopped being consumed deliberately at all, and whole categories of life — the night shift, the evening class, the lit street — appeared as if for free. Calculation traced the same arc: "computer" began as a job title, and when arithmetic became effectively free, economies did not retire the clerks so much as invent computation-dense industries no clerk could have staffed. The lessons carry over: the price collapse is half the story, and the reorganization around the newly cheap input — where nearly all the value lives — is never automatic.

EXHIBIT 2.1 • STYLIZED

The utility curve: cost of a unit of cognitive work, against a flat human benchmark



Stylized rendering of the volume's premise, not a price series: the cost of a bounded, competent cognitive task (a first draft, a document review) performed by frontier systems, against a human-expert benchmark whose cost stays roughly flat. Solid segments reflect the qualitative record; dashed segments are momentum, not forecast, and no absolute magnitude is implied.

Cognition-as-utility means a unit of competent cognitive work — a draft, a review, a diagnosis-shaped opinion, a plan — acquires the cost profile of electricity: metered, ambient, and cheap relative to the human benchmark it once required. The first-order effect is not that existing work gets done cheaper. It is that **expertise loses its queue**: waiting — for the specialist's slot, the lawyer's review, the analyst's memo — has been the tax cognitive scarcity levied on everyone, and abundance refunds it. The second-order effect: thoroughness stops being a luxury good. Things nobody commissioned because cognition was dear become affordable by default:

- the second opinion on every consequential decision, not the contested few;
- the exhaustive review — every clause, every anomaly — instead of the sample;
- the lesson paced to one learner rather than the median of thirty;
- the public service that answers the same day it is asked.

This volume rests on one distinction: **abundance of capability is not abundance of benefit**. Capability is declared upstream, in a handful of laboratories; benefit is collected downstream, across millions of firms, agencies, classrooms, and clinics. Between the two stand three mechanisms — **deployment** (systems doing real work in real workflows), **diffusion** (adoption spreading past the sophisticated few), and **trust** (permission to enter the domains where the stakes, and therefore the value, are highest). They are this report's spine, and the reason for its title.

THE CLAIM IN ONE SENTENCE

Abundance is not what the models can do; it is what societies actually collect — and collection runs through deployment, diffusion, and trust. A country can host frontier laboratories and still miss the dividend. A country with none can still collect it. The difference is built, not inherited.

02

Five abundance frontiers

Abundance talk fails when it stays abstract — "growth", "productivity", "transformation". This chapter names five domains where cheap cognition converts into things people recognizably want, and describes what abundant actually looks like in each.

2.1 Science, accelerated

The research loop — read, hypothesize, design, run, interpret, publish — is cognitively bottlenecked at every joint. The literature outgrew human reading long ago; experiment design leans on the intuition of a few; interpretation queues behind scarce senior attention. Abundant cognition attacks each joint: synthesis across an entire field on demand, hypotheses generated and triaged in volume, simulation cheap enough to precede synthesis, replication routine because its cost collapsed. The prize is not any single discovery but a faster loop — and because science is upstream of medicine, energy, and materials, this frontier pays a meta-dividend: it compounds every other frontier on this list.

2.2 Medicine without waiting rooms

Health systems ration by clinical attention: the scarce resource is an expert hour. Abundance looks concrete here — documentation handled ambiently so clinicians face patients instead of screens; triage that never sleeps and never resents the third question; a competent second opinion as the default for every consequential diagnosis, not the privilege of the well-connected; discovery loops for molecules and trials compressed from years toward months. Medicine is also the frontier where the permission-slip logic of Chapter 04 binds hardest: nothing here diffuses on capability alone, and nothing should.

2.3 A tutor for every learner

One of education research's oldest and most stubborn findings is that individual tutoring dramatically outperforms classroom instruction — and that no society has ever afforded it at scale. Abundant cognition makes the tutor a commodity: infinitely patient, paced to mastery rather than to the median of thirty, available in every language and at every

hour. The realistic version keeps teachers at the center — redeployed from content delivery toward motivation, judgment, and care — and treats the tutor as the first mass-produced complement to human teaching rather than its replacement.


2.4 Energy, materials, and the physical dividend

Cognition feeds back into atoms along two channels. The direct channel is discovery — candidate materials, catalysts, and storage chemistries screened in silico at a pace no laboratory could match, with physical experiments reserved for the shortlist. The indirect channel is throughput: grids balanced better, industrial processes tuned continuously, and the vast cognitive overhead of building things — permitting, siting, compliance paperwork — compressed from years toward weeks. The loop closes on itself: cheaper energy makes cheaper compute makes cheaper cognition makes better energy technology. Abundance, once started, feeds abundance.

2.5 State capacity, the quiet frontier

The least glamorous frontier may be the most consequential for how the decade feels. Court backlogs, visa queues, benefits adjudication, procurement cycles — these are cognition queues, and citizens experience them as the state's indifference. Abundance looks like a government that answers: applications resolved in days with reasons attached, regulation navigable in plain language, casework that no longer punishes the poor with waiting. The dividend here is paid in legitimacy as much as output — and, as Chapter 03 shows, this is precisely the frontier currently adopting slowest.

EXHIBIT 2.2 • FAIR LABS ASSESSMENT (ILLUSTRATIVE INDEX)

The mismatch: latent dividend vs readiness to absorb, five frontiers against today's adoption pools

FAIR Labs analytical construct; positions express the report's qualitative judgment on 0–10 scales, not measurement. Gold markers are the five frontiers; gray markers are reference domains where adoption currently concentrates. The pattern to notice: adoption pools where stakes are low, while the dividend waits where readiness — skills, integration, assurance — has not yet been built.

Read the exhibit cynically and it says the market is failing. Read it correctly and it says the market is *rational*: firms deploy where errors are cheap and integration is easy. Moving the mass of adoption up and to the left — into the trust-gated, high-dividend quadrant — is not something capability improvements do on their own. It is what the rest of this report is about.

03

The diffusion problem

The economic history of every general-purpose technology says the same impolite thing: the invention is not the value. Electricity rewarded the factories that redesigned themselves around it, decades after the dynamo worked. AI is running the same script, faster — and the bottleneck has already moved from the models to the adopters.

Adoption is not one act but three margins, and the dividend requires all of them. FAIR Labs tracks diffusion as **breadth** (how many organizations use capable systems at all), **depth** (how far into the workflow they reach — from autocomplete at the edges to redesigned processes at the core), and **height** (how high the stakes of the tasks entrusted). The 2026 pattern is consistent across sectors: breadth is impressive, depth is thin, height is low. Most organizations use AI the way early factories used electric motors — bolted onto the old architecture, replacing the steam engine but not the layout — and early-factory results follow: real but modest gains, concentrated in the tasks that changed least.


Depth is where the compounding lives, and depth is expensive. It requires **process re-design** (the workflow rebuilt around cheap cognition, not the model bolted onto the old one), **skills** (using a model is easy; redesigning a claims department around one is a scarce discipline), **data plumbing** (systems can only be delegated work they can see), and **trust** (monitoring, fallbacks, and accountability that make delegation defensible to a board, a regulator, or a front page). Note what that list is: the complements are mostly organizational, not technical. This is why capability gains keep outrunning realized value — every complement on the list moves at the speed of institutions.

The lag is not evenly distributed. Large sophisticated firms build the complements in-house. **Small and mid-sized enterprises** — most employment in most economies — cannot: no integration engineers, no process designers, no risk staff to satisfy, and margins too thin to fund experiments that might fail. The **public sector** faces a crueler asymmetry: the upside of adoption is diffuse and slow, while the downside is a named official in a headline. Both gaps are structural, which is the good news — structure responds to

policy. Agricultural economies once solved exactly this problem with extension services that walked new methods onto every farm; Chapter 07 proposes the cognitive equivalent.

EXHIBIT 2.3 • SCENARIO CONSTRUCT

Same capability, three futures: diffusion into high-stakes work under three trust trajectories



FAIR Labs scenario construct, not a forecast: the share of high-stakes cognitive work delegated to AI systems under three assumptions about assurance and trust, holding capability identical across scenarios. The shaded region — the gap between assured and stalled diffusion — is the volume's central quantity: value that exists in every scenario's models and is collected in only one of them.

Exhibit 2.3 makes the volume's pivot explicit. The three curves assume the *same models*. What differs is whether delegation into consequential work is made defensible — evaluated, monitored, insurable, explainable to the people it affects. That variable is not set in the laboratory. It is set by the assurance work of Volume 01 and the governance work of Volume 04, which is why the next chapter treats trust not as a virtue but as an input.

04

The trust dividend

Volume 01 measured the capability–assurance gap in safety terms. This chapter measures it in money. Every year the gap stays open, the highest-value domains keep their doors closed — and the abundance dividend quietly forfeits its largest line items.

There are, economically speaking, two AI economies. In the **low-stakes economy** — content, code suggestions, search, back-office drafting — errors are cheap, reversal is easy, and adoption needs no one's permission. That economy is real and growing, and it is where nearly all of today's diffusion pools. The **high-stakes economy** — medicine, energy, courts, credit, infrastructure, government adjudication — is where most of the latent dividend sits, and it does not run on capability. It runs on *permission*: of regulators who must license, insurers who must price, procurement officers who must sign, boards who must answer for failure, and publics who must consent to being on the receiving end. Every one of those permissions is, at bottom, a demand for assurance.

This is why safety and assurance are inputs to growth, not taxes on it. Economies have seen the pattern before: audited accounts did not slow capital markets, they created them, by turning claims strangers make about themselves into things strangers can rely on. Aviation did not grow despite its certification and crash-investigation machinery; that machinery is what turned a daredevil technology into infrastructure the public boards without thinking. Evaluations, monitoring, incident reporting, independent verification — the apparatus of Volumes 01 and 04 — play the same role for AI: they are how capability clears customs into the domains where it is worth the most.

The dividend is paid in proportion to trust. Capability sets the ceiling; assurance sets the schedule — and a society that treats safety as a brake will be out-collected, domain by domain, by one that treats it as a road.

THE CENTRAL ARGUMENT OF THIS REPORT

The counterfactual matters, because untrusted capability does not politely wait. It goes **shallow** — confined to the low-stakes economy, where returns thin as everyone arrives. And it goes **shadow** — used unofficially inside consequential workflows, unmonitored and undeclared, importing exactly the risks assurance exists to manage while generating none of the trust that would widen legitimate use. Worse, trust behaves like a common-pool resource: a single spectacular failure in one deployment reprices diffusion everywhere, shifting the entire economy from the upper curve of Exhibit 2.3 toward the lower one. Underinvesting in assurance is thus not a private economy; it is a levy on everyone else's curve.

5

frontiers where the latent dividend concentrates — science, medicine, education, energy, state capacity

4/5

frontiers gated primarily by trust and integration rather than by model capability (FAIR assessment)

3

margins of diffusion — breadth, depth, height — all of which trust must underwrite

THE PERMISSION SLIP

The fastest route to a large AI economy runs through the unglamorous machinery of Volume 01: evaluations that bind, incidents that get investigated, auditors who can look, logs that cannot be quietly edited. Not because virtue demands it — because that is what insurers price, what regulators accept, what boards approve, and what publics will forgive. Assurance is what abundance looks like from inside a procurement office.

05


The distribution question

A dividend can be collected and still be pocketed narrowly. Who ends up holding the gains from cheap cognition is not settled by the technology — it is settled by market structure, bargaining power, and policy, which is to say: it is still open.

The gains flow through three default channels. **Capital** collects through the owners of models, compute, and the equity of adopting firms. **Labor** collects where human skills complement the systems — judgment, relationships, accountability — and command a wage premium for it. **Consumers and the public** collect when competition forces gains through to prices, quality, and services. The historical record is blunt about the deciding variable: where downstream markets stayed competitive, technological windfalls passed through to buyers within a generation; where a bottleneck could be held — a network, a standard, a right-of-way — rents persisted for as long as the bottleneck did.

AI's economics tilt the defaults toward concentration, through four ratchets. The **capital intensity** of the frontier limits who can play at all. **Data feedback loops** let deployed systems improve with use, compounding incumbency. **Talent gravity** pulls the scarcest builders toward the deepest pockets. And **distribution moats** let a handful of platforms decide what the default assistant is for billions of people. Each turn of each ratchet is individually defensible; jointly they narrow the set of collectors year by year — the market-concentration territory mapped in Volume 03, with the compute-layer remedies held for Volume 07.

EXHIBIT 2.4 • ILLUSTRATIVE

Two ways the ledger can settle: who collects the dividend, default vs broad-collection trajectory

Illustrative shares constructed to make the volume's argument discussable — not estimates, not forecasts. "Default" extrapolates current market structure; "broad collection" assumes the levers of Exhibit 2.5 are pulled. The argument is directional: policy moves the split materially without shrinking the total.

The essential point of Exhibit 2.4 is that the two bars are the *same size*. Distribution policy done well operates on collection, not on capability: it widens who the gains reach without throttling the engine that generates them. The ledger below pairs each collection mechanism with its widening lever — and none of the levers is a cap on what models may learn to do. Confusing the two questions — how fast the frontier should move (Volumes 01 and 04) and who collects what it yields (this chapter) — produces policy that is bad at both.

EXHIBIT 2.5 • CONCEPTUAL FRAMEWORK

The dividend ledger: collection mechanisms, default collectors, and the levers that widen them

MECHANISM	COLLECTS BY DEFAULT	CONCENTRATION PRESSURE	WIDENING LEVER
Firm-level productivity	Shareholders of sophisticated adopters	High	Diffusion services and shared integration capacity for SMEs
Price pass-through	Consumers — only in competitive markets	High	Competition enforcement at model, cloud, and distribution layers
Complement wages	Workers already near the technology	Medium	Mid-career reskilling entitlements; adoption skills as public curriculum
Scientific spillovers	Everyone — slowly, and unevenly	Medium	Open publication norms; public research compute (Volume 05)
Public-service capacity	Taxpayers and service users	Low	Outcome-based procurement; state adoption programs (AB-02)
Equity appreciation	Owners of models and compute	High	Broad ownership vehicles — pension and sovereign exposure, employee equity
Infrastructure rents	Platform and standard owners	High	Interoperability, portability, and access mandates

FAIR Labs framework. "Concentration pressure" is a qualitative judgment of how strongly each mechanism drifts toward narrow collection absent policy; the levers act on that drift, not on frontier capability itself.

One boundary claim, to keep this chapter honest: no single instrument — including the universal transfers that dominate the popular debate — is the answer everywhere. The right mix is jurisdictional, and the labor-market half of the question is deep enough that this series gives it a volume of its own. Volume 08 takes up work, wages, and transition; this chapter's narrower point stands regardless: the split is a choice, and the time to choose is while the ratchets are young.

06

The failure modes of plenty

Abundance has a pathology chart of its own. None of the five failure modes below argues for slowing down; every one argues for steering — because each is cheap to prevent early and brutal to unwind late.

6.1 Dependency without resilience

Utilities earn dependence, and dependence changes the cost of failure. An economy that has rebuilt its workflows around cheap cognition — supplied by a handful of shared models, clouds, and chips — has concentrated its continuity risk into a monoculture: outages, defects, and attacks propagate to everyone at once, exactly the correlated-failure dynamic Volume 06 dissects. The lesson is not to refuse dependence; nobody un-depends on electricity. It is that fallbacks, degradation modes, and substitutability are part of the dividend, not a tax on it — abundance you cannot lean on is not abundance.

6.2 Hollowed institutions

When the cheap competent answer is always available, institutions can quietly stop maintaining the expensive capacity to produce their own. The failure arrives a decade later: the profession that must catch the system's errors stopped training the juniors who were supposed to become its seniors; the agency that outsourced its judgment can no longer evaluate what it buys. Abundance done well *redeploys* human expertise up the stack. Abundance done lazily dissolves it — and with it, the supervision the whole arrangement depends on.

6.3 Rebound and noise

Cheaper inputs get consumed more — that is partly the point, and history's efficiency gains have routinely been swallowed by expanded use. The benign version is more science, more teaching, more service. The malign version is congestion: synthetic content, automated persuasion, and machine-generated noise flooding the channels humans use to allocate attention and trust. Abundance of output is not abundance of value; when the cost of producing claims collapses faster than the capacity to verify them, the scarce factor becomes the signal — the epistemic territory Volumes 03 and 09 patrol.

6.4 Cost-disease inversion

As AI-touched sectors deflate, the household budget tilts toward what stays human-bound and scarce — housing, elder care, the artisanal and the positional. Measured inflation becomes two economies in one index: falling prices for cognition-dense goods, rising relative prices for everything cognition cannot touch. The politics of abundance will be decided there: a decade of extraordinary productivity can still *feel* like squeeze if the gains deflate the spreadsheet while the rent absorbs the raise. Central banks and statistical agencies are on this chart's critical path — hence AB-06.

6.5 Winner-take-most geography

Compute clusters where power and capital are; talent clusters where talent already is. Left alone, the production side of the dividend agglomerates into a few metros and member states while everywhere else merely consumes the utility — collecting some surplus, holding none of the equity, wages, or institutional capacity. A dividend can be national in the accounts and local in fact. The correction is unfashionable and proven: put the diffusion services, the public compute, and the procurement demand where the recipients are, not where the suppliers already sit.

SYMPTOMS WORTH A LINE IN THE BUDGET SPEECH

Four early symptoms, each observable years before the damage: critical services with no rehearsed fallback; professions whose junior pipelines are quietly collapsing; verification capacity losing ground to generation capacity; and adoption maps that match the existing prosperity map street for street. Each has a countermeasure in Chapter 07 — and all four reappear as indicators on the closing dashboard.

07

The Abundance Compact

Seven commitments, each with an owner — because a dividend nobody is assigned to collect is a dividend that stays on the table. Together they fund the complements, keep the engine competitive, widen the ownership, and make the dependence safe to lean on.

AB-01 Stand up national diffusion services

The agricultural-extension model, rebuilt for cognition: shared integration engineering, model workflows, and adoption vouchers that carry capable systems into the small firms and agencies that cannot build the complements alone.

OWNERS: **FINANCE MINISTRIES** · **DEVELOPMENT BANKS** · DEEP DIVE: VOL 08

AB-02 Buy outcomes, not licenses

Point public procurement at verified results — cases cleared, waits shortened, error rates down — so demand pulls assured systems into the public sector and pays for evidence rather than seats.

OWNERS: **FINANCE MINISTRIES** · **HEALTH & EDUCATION MINISTRIES** · SEE ALSO: VOL 04

AB-03 Fund assurance as growth policy

Evaluation, monitoring, and verification capacity belong in the growth budget, not the compliance budget: every point of assurance opens doors capability alone cannot. The permission slip is an economic asset.

OWNERS: **FINANCE MINISTRIES** · **FRONTIER LABS** · DEEP DIVE: VOLS 01, 04

AB-04 Police the ratchets while they are young

Interoperability, portability, switching rights, and access terms for essential inputs — enforced at the model, cloud, and distribution layers now, before today's advantages harden into next decade's tollbooths.

OWNERS: **COMPETITION AUTHORITIES** · SEE ALSO: VOLS 03, 07

AB-05 Widen the ownership base

Pension and sovereign-fund exposure to AI upside, employee equity in adopting firms, and philanthropic underwriting of public-good deployments — so the capital channel pays out beyond the capital.

OWNERS: **INVESTORS** · **PENSION FUNDS** · **PHILANTHROPIES** · DEEP DIVE: VOL 08

AB-06 Re-instrument the economic dashboard

Task-level price indices, diffusion depth and height measures, quality-adjusted public-service output, and inflation statistics that see both economies of \$6.4. What is not measured will not be distributed.

OWNERS: **CENTRAL BANKS** · **STATISTICAL AGENCIES** · SEE ALSO: VOL 08

AB-07 Make the dependence fail-safe

Resilience standards for utility-grade cognition — rehearsed fallbacks, graceful degradation, substitutability across providers — so the economy can lean its weight on abundance without betting its continuity on any single supplier.

OWNERS: **FRONTIER LABS** · **SECTOR REGULATORS** · **INSURERS** · DEEP DIVE: VOL 06

The compact's through-line matches Volume 01's: nothing here requires a breakthrough, a treaty, or anyone's altruism. It requires budget lines, procurement language, statistical programs, and competition cases — the ordinary machinery by which societies decide, without ever quite announcing it, who their technologies are for. The dashboard below is how to check the machinery is working; FAIR Labs will re-read every indicator in the 2027 edition.

EXHIBIT 2.6 • CONCEPTUAL FRAMEWORK

The abundance dashboard: eight indicators that reveal whether the dividend is being collected

INDICATOR	WHY IT MATTERS	HEALTHY DIRECTION
Diffusion depth in SMEs	Share of small firms using AI inside core workflows, not at the edges — the breadth-to-depth conversion	Rising
High-stakes share of AI task mix	Whether trust is actually granting permission into the domains where the dividend is largest	Rising, with assurance
Pass-through in AI-exposed services	Whether cost declines reach prices and access, or stop at margins — the consumer channel's pulse	Falling real prices
Public-private adoption gap	Whether the state is collecting its share of the dividend or falling behind the economy it governs	Narrowing
Concentration at model, cloud & distribution layers	The ratchets of \$5 in one number; the competition channel's early warning	Flat to falling
Assurance share of AI budgets	Whether the permission slip is being funded like the growth input it is	Rising
Fallback readiness in critical deployments	Whether dependence is being matched by resilience — §6.1's failure mode, inverted	Rising
Complement-wage breadth	How widely the labor channel pays — premium spreading across occupations rather than pooling in a few	Broadening

FAIR Labs framework: indicators and healthy directions are analytical guidance for tracking, not current measurements. Several correspond to statistics that do not yet exist — which is precisely the point of AB-06.

WHERE TO GO NEXT

Economists and strategists should continue to Volume 08 (work and wages) and Volume 09 (the trust infrastructure the dividend runs on). Policymakers: Volumes 04 and 07. Risk owners and insurers: Volumes 03 and 06. Boards: Volume 10. And every claim about the upside in this volume assumes the assurance agenda of Volume 01 gets built — the dividend is paid in proportion to trust.



Fair Artificial Intelligence Research Labs (FAIR Labs) is an independent research institute working to ensure that advanced artificial intelligence is developed safely, governed wisely, and shared fairly. Through rigorous technical research, policy analysis, and public engagement, FAIR Labs works with scientists, governments, companies, and communities to widen the circle of people who benefit from — and have a say in — the most consequential technology of our time.

The **Frontier Series** is FAIR Labs' flagship collection of stakeholder briefings on advanced AI — one shared evidence base, ten vantage points. Each volume stands alone; together they form a common map for scientists, executives, policymakers, workers, and citizens navigating the same transition.

VOL 01 · The Safety Horizon

VOL 02 · The Abundance Dividend

VOL 03 · Atlas of AI Risk

VOL 04 · Governing the Frontier

VOL 05 · The Alignment Agenda

VOL 06 · When Systems Can't Fail

VOL 07 · The Compute Compact

VOL 08 · Work in the Age of Cognition

VOL 09 · The Trust Infrastructure

VOL 10 · The Director's Dilemma