

Using AI in Scenario Planning: Letting It Rip or Doing The Right Thing?

Ewan Simpson

Narxoz Business School

October 22, 2025

Abstract

This paper introduces Grey Swan / Archimedes, a governed, AI-assisted foresight method that produces quarterly, evidence-linked probabilities for two postures: Do The Right Thing and Let It Rip. Public indicators drive bounded moves under persistence, corroboration, diffusion, and symmetric reversion. Evidence tiers: a canon (Tier 0), slow anchors (Tier 1), a frozen set of signals refreshed each run (Tier 2), and a context layer that does not move odds (Tier 3). Levers span energy, health, education, trust, coordination, and compute, with climate treated cross-cutting. Results are global. The architecture scales to sectors and countries and iterates. The forecasts are sobering.

Keywords

AI governance; scenario planning; foresight; evidence-based policy; strategic decision-making

1. Introduction

This paper presents the Grey Swan / Archimedes foresight model as a governed, AI-assisted foresight framework that produces quarterly, evidence-linked probabilities for two policy postures, Do The Right Thing (DTR) and Let It Rip (LIR). The framework is built for users who must justify decisions under uncertainty and who need updates that are transparent, comparable, and reproducible. Its contribution is practical: it turns public indicators into disciplined, auditable odds and pairs them with concise explanations and lived-context narratives that teams can communicate.

Evidence is organized in tiers. Tier 0 encodes a versioned canon of first principles. Tier 1 provides slow anchors that set denominators and bounds. Tier 2 is a set of public, measurable, persistent operational signals whose vintages are refreshed each run. Tier 3 is a focused application layer that produces context packs, operational vignettes, a frictions log, and candidate indicator cards

for countries, sectors, and cities; it never alters odds. Six 'Archimedes Levers' link evidence to outcomes in energy and grids, health security and workforce, education and skills, trust and information integrity, treaty-lite coordination and supply security, and compute and AI access. Climate change is treated as a cross-cutting driver that appears in Tier 1 baselines and in Tier 2 effects within these levers. Governance enforces small movement limits, requires persistence and corroboration across levers, applies a diffusion check so gains reach ordinary households, and uses symmetric reversion when signals fade. No private telemetry is used; manifests are versioned; every change is logged.

Our primary interest in this paper is the method. The results presented are global-level demonstrations of the governance and mechanics rather than country verdicts. Future applications will use the same architecture to generate sector, country, and regional forecasts, as well as institution-specific views. Because AI removes much of the friction in data collection, validation, and narrative production, the model will iterate continuously with versioned manifests, regular recalibration, ablation and sensitivity checks, and audit-ready change logs.

1.1 A Cautionary Note: AI is Not a Remedy For Structural Fractures

This paper is written in a world already marked by deep division and unequal endowments. Income and wealth gaps have widened, racism and nativism persist, trust in public and private institutions is fragile, civic knowledge is thinning, education quality is uneven, and large parts of the Global South lack basic infrastructure. In such conditions, artificial intelligence is not a magic bullet. It is, rather, a powerful set of tools whose effects depend on social choices and material constraints.

Without guardrails, AI is more likely to amplify existing divides than to close them. Capability and value concentrate around scarce inputs such as high-end computing, specialized talent, proprietary data, and reliable electricity. These inputs are unevenly distributed across regions and firms, creating cumulative advantage for early leaders. AI also reduces the cost of persuasion and deception faster than societies can build civic immunity. If provenance is not on by default in high-risk domains and if high-impact systems are not audited before deployment, information disorder is likely to grow and trust decline. In labor markets, AI adoption without large investments in training and honest assessment tends to widen polarization: high-skill workers gain, routine roles compress, and entry paths narrow. Where grids are tight, connectivity weak, or devices

unaffordable, access to AI is rhetorical rather than real. If leaders chase single metrics, they invite gaming and miss distributional effects.

The stance in this paper is therefore cautious and practical. AI is used to orchestrate a disciplined method, not to promise a cure. We bind first principles to rules, tie rules to public measures, and move judgments only when improvements are broad, persistent, and shared. The purpose is to make learning steadier and action more accountable in a world where technology on its own cannot repair structural fractures.

1.2. Artificial Intelligence and Scenario Planning

Artificial intelligence has moved from a specialized research program to a general-purpose capability that increasingly structures work, learning, and governance. Foundational advances in representation learning and attention, together with the consolidation of “foundation models,” explain much of this shift (LeCun, Bengio, & Hinton, 2015; Vaswani et al., 2017; Bommasani et al., 2021). Performance gains are not merely algorithmic; they emerge from a coupled system of algorithms, compute, and data in which scaling relationships are empirically regular. Accuracy and capability tend to improve predictably with increases in parameters, training tokens, and compute (Kaplan et al., 2020; Sevilla et al., 2022). Recent assessments indicate that industry now leads frontier development, with continued growth in training runs, model size, and energy demand, marking AI’s transition toward an infrastructural technology shaped by hardware supply, power availability, and data center build-out (Stanford HAI, 2025; Benaich & Hogarth, 2025). At the organizational level, adoption has deepened from pilots to “rewiring” of workflows, risk controls, and talent models, even as firms encounter persistent reliability and governance challenges (McKinsey & Company, 2025; Wharton Human-AI Research and GBK Collective, 2025). AI now functions as infrastructure. It accelerates analysis and coordination, but it also creates error modes and dependencies that outpace institutional oversight. Forecasting research underscores the need for cautious, auditable updating under such conditions, with explicit priors, small reversible steps, and clear reversion when signals fade (Agrawal, Gans, & Goldfarb, 2018; Makridakis, Spiliotis, & Assimakopoulos, 2018; Russell & Norvig, 2021).

The expansion of AI capability has revived interest in methods that help decision-makers act under deep uncertainty. Scenario planning is one of the twentieth century’s most enduring such methods. Originating in military and energy contexts and popularized in corporate strategy by

Shell's work in the 1970s and 1980s, scenarios were designed to widen leaders' field of view beyond point forecasts and to legitimize discussion of uncomfortable risks and strategic bets (Schwartz, 1996; Wack, 1985). Subsequent scholarship systematized the craft, clarifying steps for identifying drivers, constructing internally coherent narratives and using "signposts" to track unfolding change (Schoemaker, 1995; van der Heijden, 1996). Reviews of practice highlight the method's utility for strategic conversation, shared mental models and option discovery across public and private settings (Bishop et al., 2007; Bradfield et al., 2005; Chermack, 2011; Varum & Melo, 2010).

At the same time, the classic scenario toolkit exhibits well-documented weaknesses when the environment is volatile. Scenarios are often built in workshops by expert teams, refreshed irregularly and justified through qualitative argument. Updating tends to follow budgeting cycles rather than evidence; the link from new signals to revised judgments is opaque; and audit trails are thin, which impedes learning across cycles (Bradfield et al., 2005; Bishop et al., 2007). When conditions shift rapidly, as they are in current AI-shaped contexts, static narratives lag reality, and the social value of scenarios degrades. Moreover, even when "signposts" are specified, organizations frequently lack disciplined protocols for converting new readings into explicit probability adjustments, for reversing prior claims when signals fade or for limiting over-reaction to noisy observations (Makridakis et al., 2018; Tetlock and Gardner, 2015).

Modern AI does not predict the future. In our work, it orchestrates a method that enforces evidence and audit. Large models are well-suited to procedural consistency, enforcing rules for evidence quality, applying movement caps, documenting sources, checking arithmetic and producing structured outputs under constraints. When embedded within a governed workflow, AI becomes a tool for disciplined updating rather than a black-box oracle (Agrawal et al., 2018; Bommasani et al., 2021). This enables three improvements. First, calibration as evidence is incorporated in small, reversible steps, with explicit persistence tests and reversion rules inspired by best practice in forecasting (Tetlock & Gardner, 2015). Second, accountability, as every change is logged with source, indicator, magnitude and rationale, enabling ex post audit and teaching value (Makridakis et al., 2018). Third, separation of concerns, as exploratory reading can be kept distinct from the subset of public, replicable indicators allowed to move probabilities, reducing contamination from hype or vendor claims (Russell & Norvig, 2021).

These improvements open the door to a new era of assessing futures. In this era, scenario work is not a one-off artifact but a governed process that advances on a regular cadence, is transparent enough to be audited and is conservative in the face of noise. The point is not to replace human judgment. Rather, it is to bind judgment to clear procedures so that organisations can learn across cycles, avoid lurches driven by single studies or headlines and build legitimacy with stakeholders who rightly demand to see both the backbone of assumptions and the provenance of updates (Schoemaker, 1995; Chermack, 2011). AI's comparative advantage of fast synthesis under constraints can be directed toward method, making high-quality scenario practice feasible at the speed of contemporary change.

To locate our contribution within established futures practice, we recognize a wide toolkit already in use by governments, firms and civil society. Core "big-tent" guidance frames horizon scanning and structured foresight as routine inputs to policy and strategy (UK Government Office for Science, 2024, 2025; OECD, 2024). Within that umbrella, Delphi methods synthesize expert judgment under uncertainty, with recent methodological reviews clarifying where consensus definitions and reporting standards still vary (Schifano and Niederberger, 2025; Hasson, Keeney and McKenna, 2025). Structured design logics such as morphological analysis and cross-impact balances formalize internal consistency and interaction effects among drivers (Ritchey, 2022; Weimer-Jehle, 2006), while participatory variants combine these logics to surface stakeholder knowledge in energy and transition contexts (Kurniawan et al., 2022). Roadmapping links long-term intent to nearer-term capability and investment, in both enterprise and sectoral settings (Phaal, 2024; Spaltini et al., 2024). Normative planning traditions such as backcasting remain important where desired end-states are clear but feasible pathways are contested (van der Voorn et al., 2023; Wübbelmann and Kabisch, 2025). For stress-testing and adversarial thinking, wargaming and red-teaming provide procedural discipline and sponsor-facing guidance (NATO Allied Command Transformation, 2023; UK Ministry of Defence, 2023). Finally, dynamic modeling and decision analysis extend scenario work into policy design and timing choices, from agent-based and system-dynamics comparisons to real-options valuation under deep uncertainty (Elsenbroich, 2023; Schünemann et al., 2024; Farmer, 2025; Martello et al., 2024; Goruppec et al., 2022). We use scenario planning as the spine; other methods act as feeders and stressors (Figure 1).

Figure 1. Futures Methods

Method	Purpose	Strength	Representative recent sources
Horizon scanning & structured foresight	Systematically track emerging issues, weak signals and trends to inform strategy and policy	Portfolio-wide awareness; early signal detection; agenda-setting	UK Government Office for Science (2024, 2025); OECD (2024)
Delphi (classical, modified)	Iterative expert elicitation to structure and converge judgment on complex questions	When evidence is sparse and dispersed across experts; indicator or standard setting	Schifano & Niederberger (2025); Hasson, Keeney, & McKenna (2025)
Morphological analysis	Explore the full space of configurations of a problem by varying key dimensions	Early-stage problem structuring; option generation	Ritchey (2022)
Cross-Impact Balance (CIB) & participatory CIB	Ensure internal consistency among drivers and pathways; model qualitative interactions	Building small sets of coherent scenarios; tying stakeholder input to formal structure	Weimer-Jehle (2006); Weimer-Jehle (2023); Kurniawan et al. (2022)
Technology & strategy roadmapping	Align markets, products, technologies and capabilities over time	Translating long-term intent into staged investments and capability builds	Phaal (2024); Spaltini et al. (2024)
Backcasting	Start from a desired end-state and design plausible pathways back to the present	Normative transitions with clear goals and multiple routes	van der Voorn et al. (2023); Wübbelmann & Kabisch (2025)
Wargaming & red-teaming	Explore and test plans, assumptions and competitive dynamics	Stress-testing under conflict, contestation, or strategic interaction	NATO Allied Command Transformation (2023); UK Ministry of Defence (2023)
Dynamic simulation (SD/ABM) & model-based policy design	Represent feedbacks and heterogeneous agents to test policy options	When timing, path-dependence and behavior matter	Elsenbroich (2023); Schünemann et al. (2024); Farmer (2025)
Decision under uncertainty & real options	Value flexibility in the sequencing and timing of actions	Capital-intensive, climate-exposed, path-dependent choices	Martello et al. (2024); Gorup et al. (2022)

1.3 Why Scenario Planning Works for The Grey Swan and Archimedes Levers


Scenario planning fits this project because it combines decision relevance with intellectual honesty under deep uncertainty. We work in a world where many drivers interact, where path dependence matters, and where evidence arrives unevenly. Point forecasts and single-model simulations struggle in such settings because they require strong assumptions and can hide value judgments. Scenarios let us make those judgments explicit. Our two paths capture the key behavioral choice facing leaders, and the four outcomes convert that choice into states people can recognize. This structure is simple to communicate and to manage. It accepts that several futures are plausible, then requires public evidence to move the odds among them in small, reversible steps. It also provides a natural home for other methods. Scanning, Delphi, cross-impact work, roadmapping, wargaming and modeling feed the scenarios as context, design inputs or stress tests, but they do not by themselves determine probabilities. The scenario spine, coupled with Tier 0 guardrails and

the Wealth-Diffusion Gate, ensures that updates happen only when indicators show persistent, shared improvement. In short, scenario planning gives us a transparent container that links first principles to lived outcomes, supports quarterly learning with consistency, and keeps the focus on coordinated action rather than on speculative prediction. The next section specifies the rulebook (Tier 0), the evidence tiers (Tier 1 and Tier 2), the two paths and four outcomes, and the six levers that link public measures to lived results.

2. From Classic Scenarios To A Governed, AI-Assisted Framework

Against this backdrop, we introduce the Grey Swan/Archimedes framework summarized in Figure 2, which combines scenario planning’s conceptual strengths with an AI-assisted governance layer. The glossary below defines key terms of the model.

Figure 2. The Grey Swan and Archimedes Lever Scenario Planning Model



Glossary

Grey Swan (after Taleb’s “Black Swan”) denotes a future that sits between routine “white swans” and Taleb’s (2008) rare, high impact, Black Swans. Grey Swans are *plausible events with visible signposts in today’s data; they are often ignored because they are uncomfortable or cross-cutting, not because they are unknowable*. In short, they are foreseeable enough to prepare for and consequential enough to matter.

Archimedes Levers. A small set of practical domains where effort has outsized effect, echoing Archimedes’ principle of leverage. Move these levers together and systems shift; move them in isolation and gains dissipate.

Paths

- **Do The Right Thing** (DTR): Leaders choose material reforms before 2030 and act on public indicators with coordinated, persistent implementation.
- **Let It Rip** (LIR): Leaders defer action despite visible evidence, relying on piecemeal responses while structural frictions persist.

Outcomes

- **Full Empowerment** (FE): Capabilities are widely available, systems are resilient, and gains are broadly shared; credited only where there is multi-lever, persistent progress and an open Wealth-Diffusion Gate.
- **Managed Disruption** (MD): Shocks continue but are usually absorbed; core services hold with targeted fixes, though gaps remain.
- **Uneven Transition** (UT): Real progress in some regions and sectors alongside stagnation elsewhere; benefits are patchy and brittle.
- **Lost Control** (LC): Problems outpace institutional capacity; cascades, slow recovery, and eroding legitimacy dominate.

Archimedes Levers

- **Energy and Grids**: Generation, transmission, storage, and interconnection that deliver reliable, affordable power where it is needed.
- **Health Security and Workforce**: Surveillance, clinical capacity, and staffing that detect threats early and sustain care through peaks.
- **Education and Skills**: Teaching, assessment, and credentials that produce adaptive learners with employer-recognized proof of ability.
- **Trust and Information Integrity**: Provenance, transparency, and independent audits that make content traceable and high-impact systems accountable.
- **Treaty-Lite Coordination and Supply Security**: Narrow, enforceable cross-border compacts that keep critical inputs and data flowing and resolve disputes quickly.
- **Compute and AI Access**: Practical, safe access to capable compute and models for universities, small firms, and public-interest work, with pre-deployment safety checks.

Wealth-Diffusion Gate

- Empowerment is only credited when household financial buffers (savings) rise, profit-sharing or micro-equity reach meaningful breadth, and within-country inequality trends improve.

2.1 Overview of the Grey Swan and Archimedes Levers Model

The Grey Swan (after Taleb's Black Swan: visible, consequential risks rather than unforeseeable shocks) and Archimedes Levers (from 'give me a lever and I will move the world') model is a disciplined way to think about the future across 2030, 2040, and 2050. It does not try to guess a single future. Instead, it sets out two broad scenarios for how leaders behave and four possible outcomes inside each scenario. The model puts human agency at the center. It asks what people choose to do, what institutions make possible, and whether gains are shared. It then assigns odds to the outcomes and changes those odds only when public evidence meets clear rules. The model sits on three layers as shown in Figure 2 above.

First is Tier 0, which is the rulebook. Tier 0 comes from a named canon of ideas about people, power, learning, limits, and responsibility. Those first principles are turned into simple operational rules. Second is Tier 1, which contains slow, trustworthy sources that define the background shape of the world. Tier 1 sets ceilings, floors, and denominators and does not move the odds in any given quarter. Third is Tier 2, which contains dynamic indicators that can move the odds when they are measured in public, are persistent across readings, and pass quality checks. Tier 2 is organized around a short list of pressure points called the Archimedes Levers. The model also uses a Wealth-Diffusion Gate to ensure that the best outcome is only reached when gains are shared. Every change to the odds is small, logged, and reversible if the evidence fades. This turns scenario work from a one off exercise into a repeatable process that anyone can audit.

2.2 Tier 0

Tier 0¹ is the foundation. It is the rulebook that keeps the model honest. It is built from a canon of knowledge from antiquity to the present day, which states that evidence must be public and replicable, that base rates come before fresh news, that progress should be counted only when it lasts and spreads, and that power and technology have moral and social consequences. Tier 0 turns those ideas into working rules.

The main rules are simple. Movement is small and controlled. In any quarter, the odds for a single outcome can change by one to five percentage points, and the total change across the four outcomes in a path and a horizon cannot be more than eight points. Evidence must be public, measurable, and persistent. A single reading does not move the odds. A change must appear across at least two readings, and it should be corroborated when possible. If the signal that justified a change fades for two consecutive readings, the model reverses the change.

Evidence quality also matters. Some sources are fragile by nature, such as a working paper, a single country administrative series, or a vendor's telemetry. These can only nudge the odds by one or two points, and only if they are persistent. To have more weight, they must be replicated and defined clearly.

Full Empowerment, the best outcome, has a gate. The model requires several levers to meet their thresholds at the same time for at least two quarters, and it requires wealth diffusion. That means

¹ Appendix 1 provides more details. Full references used to develop 'Tier 0' are also listed.

household income buffers are improving, a real share of workers are included in profit sharing or micro equity pilots, and inequality trends are improving. Without this gate, no promotion to Full Empowerment is allowed.

There is also a strict rule for the scenario where leaders delay action (the 'Let It Rip' option). On that scenario, Full Empowerment is set to zero by default. Only if at least five levers meet thresholds, a compact across countries is verified, and wealth diffusion is present, and only if all of this persists, can a single percentage point be added to Full Empowerment at 2030. Even then, further movement requires more persistence.

Finally, Tier 0 fixes source governance. Tier 1 and Tier 2 are locked lists for a given version. Any change to those lists requires a new version and an explanation of what changed and why. Every accepted nudge within the model is logged with the source, the indicator, the size and direction of the change, the horizon and path, the reason for the move, any cap applied, and the rule for reversing it.

2.3 Tier 1

Tier 1 contains 'slow anchors'. These are public, widely used sources that set the background shape of the world. Examples include population projections, education indicators, global economic outlooks, energy outlooks, electricity access, debt statistics, and technical inventories such as the global list of high performance computing systems² sourced from agencies such as IEA, World Bank, OECD, IMF and UN. Tier 1 helps us say what is realistic. It sets ceilings, floors, and denominators. It tells us what share of a population is even in scope for a change, and how large a change might be in context. Climate is not a separate lever. It sets background constraints in Tier 1 through widely used sources on warming pathways, extreme heat days, water stress, and related baselines, while Tier 2 registers operational impacts within energy, health, coordination, and compute as discussed further below.

Tier 1 does not move the odds within a quarter. It is not about sudden change. It is about the context that changes slowly. When an official issuer releases a new edition, the model increments to a new version, records the change, and continues from there. This approach prevents

² To maintain flow, a listing of references is provided at the end of the paper

overreaction and ensures that the model's base rates are aligned with the most recent, reputable anchors.

2.4 Tier 2

Tier 2 is where movement comes from. It contains dynamic indicators³ that are public, measurable, and persistent. Each indicator belongs to one of the Archimedes Levers. If an indicator clears its threshold across at least two readings, and if the evidence is corroborated and not fragile, the model may nudge the odds within the limits set by Tier 0. If the indicator is policy fragile, the model can still respond, but the move is capped at one or two points. If the signal fades, the move is reversed.

Tier 2 indicators are chosen because they link directly to lived outcomes. They track whether energy can be delivered when and where it is needed, whether health systems can see trouble early and keep staff in place, whether teachers have the capacity and assessments reflect real learning in an age of AI, whether content is traceable and high impact systems are audited before use, whether countries can cooperate in simple and fast ways, and whether compute is open enough for small firms and public interest projects to make use of it. Tier 2 also carries the Wealth-Diffusion Gate. This draws on estimates of household savings/buffers, data on share ownership distribution and income and inequality from both international (world Bank Gini-coefficient data) and national statistical sources. Without shared gains, technical progress does not count as empowerment.

2.5 Tier 3: Interpretive and Customization Layer.

Tier 3 is an interpretive layer that instantiates the model for specific personas, sectors, and geographies. Its purpose is twofold. First, it customizes outputs so that Tier-1 anchors and Tier-2 indicators are rendered as lived experience for an intelligent generalist reader (e.g., how repeated heat-day curtailments translate into missed shifts or clinic queues). Second, it tests model completeness by surfacing plausible gaps in measurement and by proposing candidate public sources for Tier-2 consideration. Tier 3 draws only on public facts and ordinary administrative textures; it does not update odds. Any movement requires elevation to Tier 2 with a defined indicator, public provenance, and persistence across update cycles. Tier 3 outputs are

³ See footnote 2

standardized: three short operational vignettes per declared context (2030, 2040, 2050), a one-line contrast note on how circumstances change by geography or income, a simple frictions log of recurring public problems, and up to three candidate indicator cards marked Track, Pilot, or Propose. These are public-facts-only artifacts that do not move odds. Any promotion to Tier 2 requires a new version, public provenance, persistence across cycles, and non-redundancy.

2.6 Scenarios

The model runs two scenarios for how leaders behave. These are not minor variations. They reflect fundamentally different choices.

2.6.1 Scenario 1: Do the Right Thing

The first scenario is 'Do The Right Thing'. Here, governments, firms, and institutions choose material reforms before 2030. They act on what the indicators show. Progress is uneven, but the direction is constructive. Under this scenario, the model allows the odds to move toward better outcomes when several levers clear thresholds at the same time and keep doing so.

2.6.2 Scenario 2: Let It Rip

The second scenario is 'Let It Rip'. Here, leaders delay or avoid the necessary choices even though the evidence is visible. Under this scenario, the model assumes that drift and fragmentation will dominate, and that disorder will grow unless there is an unusual surge of multi lever progress that persists. Full Empowerment is set to zero by default on this scenario. Only an exceptional combination of results, including a verified compact across countries and wealth diffusion, allows a tiny improvement, and only if that improvement holds across readings.

2.7 Four Possible Outcomes – from Full Empowerment to Lost Control

'Full Empowerment' is the best outcome. It is not only technical capability. It is capability with resilience and fairness. Power is reliable enough for people and firms to plan with confidence. Health systems can manage peaks without burning out their staff. Schools and universities assess with AI built in, teachers are trained, and graduates hold credentials that employers respect. High risk content carries provenance by default, and high impact AI systems are audited before deployment with public summaries. Compacts across borders exist and work in practice, not just on paper. A share of near frontier compute is reserved and actually used for public interest work

and small firms. Household income buffers are rising, profit sharing or micro equity reaches a real share of workers, and inequality trends are improving. Full Empowerment is only reached when several levers meet their thresholds for two consecutive quarters and the Wealth-Diffusion Gate is open.

'Managed Disruption' is a world where shocks continue but the system absorbs them more often than not. Power is tight in some places and stable in others. Health systems cope better than before, but not everywhere and not always. Teachers and students use AI well in many classrooms, but coverage and quality still vary. Provenance and audits exist in the most sensitive areas, yet coverage is incomplete. There may be a compact in one domain, but not across all the places that matter. Small firms have better access to compute, but queues and costs still bite. Gains are real, but they are not secure.

'Uneven Transition' is the most common shape of a mixed future. Some regions and sectors move ahead while others lag. Some grids add capacity and reduce curtailment, while others are stuck in queues. Some education systems move to AI integrated assessment and teacher training, while others ban the tools or turn a blind eye. Some sectors build credible provenance and audit regimes, while others do not. Coordination across borders is fragile. Compute access and skills gather in a small number of firms and locations. Wealth gains cluster, and political risk remains high. It is a patchwork. People in the leading regions do well. Others feel left behind.

'Lost Control' is not a dramatic collapse in a single day. It is a state where problems outrun the capacity to manage them. Power shortages recur and hold back investment. Health systems lose staff and fail during peaks. In education, assessment integrity erodes because the old tools no longer work and new tools are not in place. Mis- and dis-information outpace provenance and audit. Cross border coordination fails when it is most needed. Compute access concentrates further. Household income buffers reduce and inequality worsens. Recovery from shocks is slow and expensive. On the scenario where leaders delay action, this outcome dominates across time.

2.8 The Archimedes Levers

The Archimedes Levers (shown in Figure 2) are the small set of pressure points that move the whole system when they shift together. They are chosen because they multiply one another, and because progress can be measured in public. Climate pressures are handled inside the levers rather than as a seventh lever. They show up in practical measures such as grid strain during heat,

hydropower variability, health peaks from heat and vector shifts, learning disruption, and cross-border compacts for water and inputs. This keeps measurement practical and avoids double counting.

The first lever is energy and grids. Modern life depends on clean, reliable power, and data centers increase demand. The model looks for material and sustained reductions in connection queues and curtailment across peak seasons, and sufficient storage and interconnection to meet projected load growth. Meeting these thresholds for two quarters is a strong sign that the rest of the system is robust.

The second lever is health security and the workforce that sustains it. The model looks for health surveillance data that covers about 80% of the population (timely syndromic feeds and genomic sequencing/wastewater coverage that together capture around 80% of people with defined turnaround times), for staffed bed ratios that are stable or improving during peaks, and for a meaningful drop in critical staff attrition. These measures show that systems can see trouble early and reduce the potential for illness.

The third lever is education and skills. The model looks for widespread AI-aware assessment with trained teachers and employer recognized micro credentials with real uptake. These measures show that people can learn and adapt at the speed the world now demands.

The fourth lever is trust and information integrity. The model looks for provenance turned on by default in high risk content and pre-deployment audits by default in high-impact systems. Detectors can be useful as advisory tools, but proof comes from process, traceability, and published audits.

The fifth lever is 'treaty lite' coordination and supply security. The model is realistic about how slow formal treaties can be. It looks for a compact across at several jurisdictions, with a clear scope, a public monitoring system, and dispute and input shock resolution times measured in days. It also looks for evidence that show the compact works in practice.

The sixth lever is compute and access. The model looks for a proportion of near frontier compute to be reserved and actually used for public interest work and small firms, for onboarding of small firms at a density of about one hundred and fifty per million people, and for safety evaluations

before deployment of frontier adjacent systems. These measures show that new capability is not hoarded and that more actors can take part.

Across all levers sits the Wealth-Diffusion Gate. The model requires a rise in the share of households with a three-month savings buffer, the spread of profit sharing or micro equity pilots to a real share of workers, and an improving inequality trend. Without this gate open, the best outcome is not available. The system might become more capable, but it would not be empowering.

This is the model in full. It begins with first principles, turns them into rules, anchors those rules in slow public facts, and allows movement only when dynamic, public indicators show persistent and shared progress. It rewards breadth and patience and it makes every change traceable.

2.9 A Step Forward in Scenario Planning

The Grey Swan and Archimedes Levers framework represents a clear step forward in automated scenario planning. Traditional scenarios are crafted in workshops, refreshed irregularly, and justified through expert narrative. They often lack a consistent update cadence and a transparent audit trail. By contrast, this model binds first principles, measurement, and governance into a single process that can be reproduced and defended.

The foundation is a canon of ideas about people, power, learning, and limits. Rather than leaving those ideas implicit, the canon is named and versioned, then translated into operational rules. This becomes Tier 0, the rulebook. Tier 0 sets small, reversible movements, insists on persistence across readings before any change is made, and forces reversion when signals fade. It defines the asymmetry between the two paths. On Do The Right Thing, progress must be broad and persistent before odds improve. On Let It Rip, Full Empowerment is zero by rule unless an exceptional, multi-lever surge is verified. Tier 0 also fixes the Wealth-Diffusion Gate. Gains are not counted as empowerment unless household buffers rise, profit sharing or micro equity reach real breadth, and inequality trends improve.

Tier 1 and Tier 2 complete the evidence pipeline. Tier 1 holds slow anchors such as population, education, energy, and macroeconomic baselines. These sources set denominators and ceilings and never move odds in quarter. Tier 2 defines dynamic, public indicators that can nudge odds within Tier 0 limits when they clear thresholds with corroboration. Examples include grid

connection medians and curtailment across peak seasons, teacher training coverage and AI-integrated assessment, provenance by default in high-risk content, verified cross-border compacts with fast dispute resolution, and reserved near-frontier compute that is actually used by small firms and public interest projects. Policy-fragile evidence is capped at one or two percentage points until it is replicated. Tier 3 allows browsing and context while keeping odds movement clean, since it cannot move numbers without promotion and a version bump.

This structure makes scenario work auditable and teachable. Every accepted nudge is logged with source, indicator, direction, size, horizon, rationale, any cap applied, and the reversion rule that will unwind it if the signal fades. The quarterly cadence and the eight-point net cap prevent lurches. The promotion pipeline from Tier 3 to Tier 2 turns exploration into method rather than anecdote. The in-chat header, operator checklist, and move-log template make the whole thing replicable by someone other than the author.

Human judgment still matters. Curating and amending the canon, setting thresholds, and handling contested contexts remain human work. The risks are well known. Measurement can be gamed, evidence can be fragile, and what is measurable can crowd out what is important. The framework mitigates this by using families of indicators, insisting on public scoreboards, capping fragile signals, and reversing moves when persistence fails.


In short, the model shifts scenario planning from expert storytelling to a governed, evidence-linked, machine-auditable process. It keeps agency and ethics explicit, treats progress as something that must be earned and shared, and shows how orchestration can deepen quality without sacrificing discipline. It is critical to note that the model uses only public sources. No private telemetry is ingested. Tier-0 canon is fixed by version; Tier-1 is frozen except when an official release appears; Tier-2 is a fixed indicator set whose vintages are refreshed on each run; Tier-3 produces narratives and experiments that do not affect odds. Each run records source and code hashes, observation periods, vintage dates, and high-level QA flags so past results can be reproduced. Every odds change is logged with posture, indicator family, direction, and the persistence, corroboration, and diffusion checks that justified it. If a signal fades on later runs, the reversion rule applies in the same step size and direction to unwind the prior move.

3. Model First Run Outcomes

3.1 Results: projected outcomes under the current specification


As noted, the framework models two behavioral paths and four possible outcomes at three time points. Figures 3 and 4 show the results based on the model as of October 30 2025. In this paper, the focus is global, since this is a foundational paper which will inform a range of other ongoing projects at different geographies and sectors. The first path, Do The Right Thing, assumes that public and private leaders choose major material reforms before 2030. The second path, Let It Rip, assumes that action is deferred despite visible evidence. Within each path, four outcomes are assessed: Full Empowerment (capability is widely available, systems are resilient, and gains are shared), Managed Disruption (shocks continue but are absorbed more often than not), Uneven Transition (progress is real but patchy across places and sectors), and Lost Control (problems outrun institutional capacity and recovery is slow).

Figure 3. Do The Right Thing



Under Do The Right Thing, the 2030 center of gravity is Uneven Transition, with a modest Full Empowerment rise by 2040–2050 conditional on breadth and diffusion. However, Lost Control rises in prominence over time.

Figure 4. Let It Rip



Note: *Full Empowerment is set to zero on this path by rule. A narrow, evidence-based exception allows only a one-point increase at 2030 if multiple strict conditions are met and sustained.

Under Let It Rip, Lost Control dominates at all horizons; Full Empowerment remains zero by rule save for the narrow 2030 exception noted above.

These figures are conservative by design. Probabilities change only when public, measurable indicators show improvement that persists across time and across several domains. In the absence of such evidence, the priors stand.

3.2 Why Do These Outcomes Dominate?

3.2.1 The Central Driver: Coordinated, Multi-Domain Progress is Improbable at Global Scale

The model links better futures to progress that is both broad and persistent. It requires simultaneous improvement across several practical domains that matter in everyday life:

- reliable and well-connected electricity systems;
- public health capacity and workforce stability;
- education that uses modern assessment with trained teachers and recognized credentials;
- trustworthy information with provenance and pre-deployment audit of high-impact systems;
- narrow but operable cross-border compacts for critical supplies;
- meaningful access to advanced computing for small firms and public-interest work.

In addition, it requires evidence that gains are shared, not merely created. Households should build financial buffers, profit sharing or micro-equity should reach a real share of workers, and inequality trends should improve.

Given current institutional incentives and frictions, the likelihood that many jurisdictions achieve simultaneous, publicly verifiable improvements across this set of domains and sustain them over multiple reporting periods is low. This is the single most important reason the Uneven Transition outcome is most likely on the reform path in 2030, and why Lost Control dominates the deferment (Let It Rip) path at every horizon.

3.2.2 Domain-specific Challenges

- **Electricity systems :Delivered power, not just generation.** New generation is not enough. Connection queues, curtailment during peak seasons, and network constraints frequently limit delivered power. Interconnectors and long-duration storage require long lead times, coordinated permitting, and capital. As a result, near-term improvements often appear as isolated gains rather than system-wide change.
- **Public health and workforce: Human capacity before tools.** Surveillance coverage and surge capacity rely on sustained investment in people and data, not tools alone. Staffed-

bed ratios during peaks are sensitive to burnout, pay, and working conditions. Many systems can purchase technology quickly but struggle to keep skilled staff and maintain performance under stress.

- **Education and skills: Honest assessment at scale.** Assessment that incorporates AI, large-scale teacher training, and employer-recognized micro-credentials demand curricular reform, investment in human capital and institutional buy-in. Adoption tends to be uneven across regions and institutions, producing a patchwork rather than national coverage.
- **Trust and information integrity: Provenance by default.** Provenance “on by default” for high-risk content and credible audit summaries for high-impact systems require shared standards and enforcement. Pilots exist, yet coverage and reliability lag. Without visible breadth, confidence remains fragile.
- **Cross-border coordination and supply security: Compact drills, not communiqués.** The framework does not rely on grand treaties. It asks for narrow compacts with fast dispute resolution and routine drills. Even this restricted ambition is difficult to institutionalize across multiple jurisdictions with different interests and legal regimes.
- **Compute and AI access: Reserved compute actually used.** Genuine access means reserved capacity that is actually used by small firms and public-interest projects, coupled with practical onboarding and safety checks prior to deployment. In practice, capacity, expertise, and attention concentrate in a small number of hubs.
- **Sharing of gains: Gains must diffuse.** Evidence that gains are broadly shared is essential. Where household buffers stagnate, profit sharing is rare, or inequality worsens, claims of empowerment are not credited, even if technical capability improves.

These factors interact. Tight grids slow digital delivery. Health stress amplifies social and economic shocks. Weak integrity systems raise political temperatures and complicate coordination. Concentrated compute and narrow wealth gains reduce support for reform. The interaction of these frictions explains why Uneven Transition dominates the reform path in the near term and why Lost Control dominates the deferment path throughout.

3.2.3 Trajectories Over Time

On the reform path, a patchwork is the most plausible near-term result. Institutions require time to convert policy into operating capability and to spread benefits beyond early adopters. The

model therefore assigns a high probability to Uneven Transition in 2030. By 2040 and 2050, the probability of Full Empowerment rises modestly, reflecting the possibility that some jurisdictions will string together sufficient multi-domain progress and demonstrate sharing of gains. The probability of Lost Control also grows by 2050, acknowledging that partial, contested, or slow reforms can allow shocks to compound.

On the deferment path, inaction and piecemeal efforts preserve system frictions. The result is high and rising probability for Lost Control. A narrow exception allows a very small improvement if many domains improve together and benefits are shared, and if this is sustained across time. The bar is intentionally high because drift is persistent and isolated improvements are easily absorbed by larger constraints.

3.2.4 Why the Update Rules are Deliberately Conservative

The framework moves probabilities in small steps, only after repeated public readings, and reverses prior moves if the evidence is not sustained. It also caps the influence of early or narrow studies until they are replicated. These rules are not administrative details. They are safeguards against overreaction to headlines and premature celebration of pilots. They protect credibility, encourage learning across reporting cycles, and create a clear paper trail that allows outsiders to understand how and why judgments changed.

3.2.5 Interpreting Each Outcome in Light of the Drivers

- **Full Empowerment.** This outcome remains rare in the near term because it requires simultaneous, sustained progress in several domains and clear evidence that benefits are broadly shared. It becomes more plausible by 2040 and 2050 in jurisdictions that align reforms and diffusion.
- **Managed Disruption.** This outcome reflects systems that cope with shocks more often than not, yet still carry gaps in coverage. It is a plausible waypoint on the reform path. Where pilots become policy and coverage widens, it can feed into Full Empowerment. Where progress stalls, it slips back toward a patchwork.
- **Uneven Transition.** This is the most likely near-term state on the reform path. It captures real gains in some places and sectors alongside stagnation elsewhere. It is also the point where many efforts stall if complementary domains do not advance.

- **Lost Control.** This outcome describes regimes where cascades outrun capacity. On the deferment path, it is the dominant state because isolated wins rarely overcome structural frictions. On the reform path, it is a persistent risk that grows if reforms remain narrow or contested.

3.2.6 What Evidence Would Change the Picture?

The projected outcomes will change when public indicators show simultaneous, sustained improvement across several domains and when there is credible evidence that gains are shared. Concretely, for energy, that would include shorter connection times and lower curtailment across peak seasons, stable performance of health systems during surges, widespread AI-aware assessment with trained teachers and recognized credentials, provenance by default in high-risk content and public audit summaries for high-impact systems, at least one narrow cross-border compact that resolves disputes and input shocks within days, reserved advanced computing that is actually used by small firms and civic projects, and rising household savings alongside broader allocation of profits or equity. When such evidence persists across multiple reporting periods, conservative but cumulative shifts toward better outcomes are warranted.

Under present conditions, the remote likelihood of coordinated, multi-domain, and persistent global action is the decisive driver of these results. The framework therefore locates the near-term (2030) center of gravity in Uneven Transition under timely reform, and in Lost Control under deferment. The purpose of stating this plainly is not to invite fatalism, but to identify the exact public signals that would justify moving the odds toward more empowering futures.

4. Conclusion

This paper has argued that scenario planning can be modernized without losing its original purpose. The tradition gave leaders a language for uncertainty and a discipline for widening their field of view. Its limits were also clear: irregular updates, weak links between evidence and revisions, and sparse audit trails. The framework presented here advances the practice by binding first principles, measurement, and governance into a single, reproducible process. It does not claim to predict the future. It claims to make learning about the future steadier, more transparent, and more useful for action.

The key advance is procedural. A named canon of ideas about people, power, learning, and limits is translated into rules that the system actually follows: small and reversible adjustments, persistence tests across readings, reversion when signals fade, and a gate that withholds the best outcome unless gains are broadly shared. Slow anchors fix the background shape of the world. Dynamic public indicators tied to concrete domains determine when small probability shifts are justified. Exploratory material is allowed for context but kept out of the mechanism that moves numbers unless it is promoted with clear criteria. Every change is logged. Every rule is applied the same way each quarter. This turns scenarios from occasional narratives into a standing method.

The second advance is the combination of orchestrated expertise and AI. Experts set the canon, choose the domains that matter, and define thresholds that link to lived outcomes. AI then enforces the method. It checks whether evidence is public and measurable. It verifies that signals persist across time. It applies movement caps and records reversion when the case weakens. It keeps logs in a standard format and produces outputs that others can audit and teach from. This division of labor matters. Human judgment remains at the center, where it belongs, but it is channeled through procedures that reduce noise, bias, and drift. The result is a system that can update at the tempo of contemporary change without lurching from headline to headline.

The results reported here at the global system level are sobering. On the optimistic path, where reforms are chosen in time, the most likely near-term state is a patchwork of progress and stagnation. On the deferment path, where action slips despite visible signals, the dominant state is one in which shocks outrun capacity and recovery is slow. These conclusions do not arise from a single dataset or from a fixed ideology. They arise from the practical difficulty of achieving simultaneous and sustained improvement across several domains that touch daily life: reliable power, staffed and visible health systems, honest assessment and teacher capacity, traceable content and public audits for important systems, narrow cross-border compacts that work in practice, and real access to advanced computing for small firms and public interest work. They also arise from the requirement that gains be shared. Without evidence that households are building buffers, that profit sharing or micro-equity is spreading, and that inequality trends are improving, we do not award the best outcome.

The implication is clear. Absent decisive coordination, AI will deepen existing divisions. Reliable electricity must be delivered, not only announced. Health systems must see trouble early and hold staffed capacity through peaks. Education systems must adopt AI-aware assessment, train

teachers at scale, and confer credentials that employers recognize. High-risk content must carry provenance by default, and high-impact systems must publish pre-deployment audit summaries. Partners must operate at least one narrow cross-border compact with dispute and input-shock resolution measured in days. Advanced computing must include reserved capacity that small firms and public-interest projects actually use, with practical onboarding and safety checks. Across all of this, gains must diffuse. Household savings should rise, profit sharing or micro-equity should reach real breadth, and inequality trends should improve.

This paper has shown how to govern scenario work so that claims move only when public indicators clear thresholds across time and across domains. The message here is direct. Only coordinated, measured, and persistent effort across the domains named above will shift the odds toward more empowering futures. If we want better odds, we must act in ways that are coordinated across domains and persistent across time. This is not a call for grand declarations. It is a call for a narrow set of measurable commitments that different jurisdictions can make and publish. Connection times and curtailment should be reported every quarter. Health systems should publish staffed capacity during peaks and attrition in critical roles. Education systems should document the share of learners assessed in AI-aware ways and the share of teachers trained to use these tools well. Media and technology firms should turn provenance on by default in high-risk categories and release public summaries of pre-deployment audits for high-impact systems. Partners should build at least one compact that can resolve disputes and input shocks within days, then drill it and show the results. Operators of advanced computing should reserve capacity that small firms and public projects can actually use and should publish utilization. Finance and labor ministries should track household savings and the breadth of profit sharing or micro-equity pilots. These are not abstract indicators. They are the signs that systems are changing in ways that matter.

We are headed, at best, into a very uncertain future. The evidence assembled here suggests that current actions are not yet commensurate with that uncertainty. The framework is deliberately conservative. It moves probabilities only in small steps and only after repeated public readings. If progress is narrow, temporary, or unequally distributed, it will not register as a change in the odds. That conservatism is not pessimism. It is a commitment to credibility. It sets a high bar for celebration and a clear path for improvement.

In sum, the path to more empowering futures runs through coordination that is narrow, practical, and measured in public, and through a discipline of updating that is both patient enough to resist noise and fast enough to learn. AI helps by making that discipline practical. Expertise helps by choosing what matters and insisting that gains be shared. Together they can move scenario planning from intermittent stories to a standing practice that earns trust and supports timely action.

AI Use Statement

This statement covers AI use in (A) writing the manuscript and (B) developing and operating the Grey Swan / Archimedes model.

A. Writing and editorial assistance

1. **AI Roles performed.** Outlining, redrafting sections, language editing, consistency checks across terminology and reference formatting.
2. **Non-roles.** AI did not originate the research question, design the governance rules, select indicator families, adjudicate results, or approve conclusions.
3. **Human control.** All passages were reviewed, edited and finalized by the author. All claims and errors are the responsibility of the author. AI is not listed as an author.

B. Model development and use

1. **Function of AI.** AI acts as an orchestrator of a governed process designed by the author. It enforces small movement limits, requires persistence and corroboration across levers, applies a diffusion check so gains reach ordinary households, and implements symmetric reversion when signals fade.
2. **Evidence architecture.** Tier 0 canon; Tier 1 slow anchors; Tier 2 a frozen set of public, measurable, persistent operational signals whose vintages are refreshed each run; Tier 3 a focused application layer that produces context packs, vignettes, a frictions log, and candidate indicator cards. Tier 3 does not move odds.
3. **Data and privacy.** Only public indicators were used. No private telemetry or personal data were ingested.
4. **Climate treatment.** Climate risk is cross-cutting. Physical and transition effects appear in Tier 1 baselines and as Tier 2 impacts inside the levers.

5. **Reproducibility.** Each run records source and code hashes, observation periods, vintage dates, and high-level QA flags. Indicator families are public; exact endpoints and transform parameters remain proprietary. Versioned manifests are available on request.
6. **Outputs and scope.** Current results are global demonstrations of the method. The architecture scales to sector, country, and regional forecasts. Because AI reduces update friction, the model will iterate with versioned releases, regular recalibration, and audit-ready change logs.

Accountability. The author accepts full responsibility for the design, operation, and interpretation of the model and for the manuscript text.

References

- Abbott, A. (1988). *The system of professions: An essay on the division of expert labor*. University of Chicago Press.
- Acemoglu, D. and Restrepo, P. (2019). Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives*, 33(2), pp.3-30.
- Acemoglu, D., Johnson, S. (2023). *Power and progress: Our thousand-year struggle over technology and prosperity*. Public Affairs.
- Agrawal, A., Gans, J., & Goldfarb, A. (2018). *Prediction machines: The simple economics of artificial intelligence*. Harvard Business Review Press.
- Ahmed, N. and Wahed, M. (2020). The de-democratization of AI: Deep learning and the compute divide in artificial intelligence research. arXiv. Available at: <https://arxiv.org/abs/2010.15581>.
- AI Action Summit (2025). *AI for the Future: Global Governance and Technological Sovereignty*. Élysée – The French Presidency. Available at: <https://www.elysee.fr/en/sommet-pour-l-action-sur-l-ia>
- Air Street Capital. (2025). *State of AI Report 2025*.
- AlixPartners. (2025). *AI in Creative Industries: Enhancing, rather than replacing, human creativity in TV and film*. <https://www.alixpartners.com/insights/102jsme/ai-in-creative-industries-enhancing-rather-than-replacing-human-creativity-in/>
- American Association of University Professors. (2025, July). *Artificial intelligence and academic professions*. AAUP. <https://www.aaup.org/sites/default/files/2025-07/TREP-Artificial-Intelligence-and-Academic-Professions.pdf>
- Aoun, J. E. (2017). *Robot-proof: Higher education in the age of artificial intelligence*. MIT Press.
- Arthur, W.B. (1989). Competing technologies, increasing returns, and lock-in by historical events. *The Economic Journal*, 99(394), pp.116-131.
- Asian Development Bank. (2025). *Asian Development Outlook (ADO) July 2025: Slower growth amid tariffs and uncertainty*. <http://dx.doi.org/10.22617/FLS250292-3>

- Atari, M., Xue, M. J., Park, P. S., Blasi, D. E., & Henrich, J. (2025). Which humans? Department of Human Evolutionary Biology, Harvard University. https://scholar.harvard.edu/sites/scholar.harvard.edu/files/henrich/files/which_humans_09222023.pdf
- Autor, D. (2015). Why are there still so many jobs? The history and future of workplace automation. *Journal of Economic Perspectives*, 29(3), 3–30. <https://doi.org/10.1257/jep.29.3.3>
- Baer, J. (2016). Domain specificity and the limits of creativity theory. *The Journal of Creative Behavior*, 50(4), 273–275. <https://doi.org/10.1002/jocb.132>
- Ballerini, V., Dominici, A., Ferracane, M. F., Menchetti, F., Noirjean, S. (2025). Stimulating creativity and grit of high school students with creative STEM activities: An RCT with noncompliance. *Quality and Quantity*, 59(1), 605–634.
- Barbot, B., Besançon, M., Lubart, T. (2015). Creative potential in educational settings: Its nature, measure, and nurture. *Education 3-13*, 43(4), 371–381. <https://doi.org/10.1080/03004279.2015.1020643>
- BCG (Boston Consulting Group). (2024). The AI Maturity Matrix. <https://web-assets.bcg.com/fe/61/6962e74b44328f148c8a9ac1002d/ai-maturity-matrix-nov-2024.pdf>
- BCG (Boston Consulting Group). (2025). The Widening AI Value Gap. <https://media-publications.bcg.com/The-Widening-AI-Value-Gap-Sept-2025.pdf>
- Bishop, P., Hines, A., & Collins, T. (2007). The current state of scenario development: An overview of techniques. *Foresight*, 9(1), 5–25. <https://doi.org/10.1108/14636680710727516>
- Bommasani, R., Hudson, D. A., Adeli, E., et al. (2021). On the opportunities and risks of foundation models. *arXiv*. <https://arxiv.org/abs/2108.07258>
- Bradfield, R., Wright, G., Burt, G., Cairns, G., & Van Der Heijden, K. (2005). The origins and evolution of scenario techniques in long range business planning. *Futures*, 37(8), 795–812. <https://doi.org/10.1016/j.futures.2005.01.003>
- Breakspear, S. (2012). The policy impact of PISA: An exploration of the normative effects of international benchmarking in school system performance. *OECD Education Working Papers*, No. 71. OECD Publishing. <https://doi.org/10.1787/5k9fdfqffr28-en>
- Brodeur, P.G., Buckley, T.A., Kanjee, Z. et al (2024) Superhuman performance of a large language model on the reasoning tasks of a physician 10.48550/arXiv.2412.10849
- Brynjolfsson, E. and McAfee, A. (2014). *The second machine age: Work, progress, and prosperity in a time of brilliant technologies*. New York: W.W. Norton & Company.
- Brynjolfsson, E., Chandar, B., Chen, R. (2025) Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence, Working Paper, August 26, 2025. Available at: <https://digitaleconomy.stanford.edu/publications/canaries-in-the-coal-mine/>
- Brynjolfsson, E., Li, Z., Raymond, L. (2023). Generative AI at Work. NBER Working Paper 31161. <https://www.nber.org/papers/w31161>
- Brynjolfsson, E., McAfee, A. (2017). *Machine, platform, crowd: Harnessing our digital future*. W. W. Norton.
- BSI (British Standards Institution). (2024). Trust in AI: A playbook for responsible adoption. <https://www.bsigroup.com/siteassets/pdf/en/insights-and-media/campaigns/trust-in-ai.pdf>
- BSI (British Standards Institution). (2025). Flourishing in the AI workforce. <https://www.bsigroup.com/siteassets/pdf/en/insights-and-media/insights/white-papers/flourishing-in-the-ai-workforce.pdf>
- Cambridge University Press & Assessment. (2024). *Preparing learners to thrive in a changing world: The view from international education*. Cambridge University Press & Assessment.
- Cambridge University Press. (2024). *Publishing futures: Working together to deliver radical change in academic publishing*. Cambridge University Press.

- Castelnovo, A. (2024). Towards responsible AI in banking: Addressing bias for fair decision-making. arXiv. Available at: <https://arxiv.org/abs/2401.08691>.
- Cazzaniga, M., Pizzinelli, C., Rockall, E., & Tavares, M. M. (2024). Exposure to artificial intelligence and occupational mobility: A cross-country analysis (IMF Working Paper No. WP/24/116). International Monetary Fund. <https://www.imf.org/en/Publications/WP/Issues/2024/06/07/Exposure-to-Artificial-Intelligence-and-Occupational-Mobility-A-Cross-Country-Analysis-549989>
- Center for AI Safety (CAIS). (2024). CAIS Impact Report 2024. Center for AI Safety. https://odw.dev/cais/CAIS_impact_report_2024.pdf
- Cerutti, E., Garcia Pascual, A., Kido, Y., Li, L., Melina, G., Tavares, M. M., & Wingender, P. (2025). The global impact of AI: Mind the gap (IMF Working Paper No. WP/25/76). International Monetary Fund. <https://www.imf.org/en/Publications/WP/Issues/2025/04/11/The-Global-Impact-of-AI-Mind-the-Gap-566129>
- Chatterji, A., Cunningham, T., Deming, D.J., Hitzig, Z., Ong, C., Shan, C.Y., Wadman, K. (2025) How People Use ChatGPT. NBER Working Paper 34255 <http://www.nber.org/papers/w34255>
- Chermack, T. J. (2011). Scenario planning in organizations: How to create, use, and assess scenarios. Berrett-Koehler.
- Cisco. (2024). AI Readiness Index.
- Coyle, D., Tan, J., Jackson, B., & Berjon, R. (2025, September). Airbus for AI: A global strategy for public value creation. Bennett School of Public Policy, University of Cambridge.
- Craft, A. (2005). Creativity in Schools: Tensions and dilemmas. Routledge.
- Davenport, T. H., & Ronanki, R. (2018). Artificial intelligence for the real world. *Harvard Business Review*, 96(1), 108–116.
- David, P.A. (1985). Clio and the economics of QWERTY. *American Economic Review*, 75(2), pp.332–337.
- De Cremer, D., Bianzino, N.M. (2023). How generative AI could disrupt creative work. *Harvard Business Review*. <https://hbsp.harvard.edu/product/H07LIA-PDF-ENG>
- DEC (Digital Education Consortium). (2025). The next era of assessment. DEC
- Digital Education Council & Pearson. (2025). The next era of assessment: A global review of AI in assessment design. Digital Education Council and Pearson.
- Ding, A. W., Li, S. (2025). Generative AI lacks the human creativity to achieve scientific discovery from scratch. *Scientific Reports*, 15, 9587.
- Dutta, S., Lanvin, B., Rivera León, L., & Wunsch-Vincent, S. (Eds.). (2024). *Global Innovation Index 2024: Unlocking the promise of social entrepreneurship* (17th ed.). World Intellectual Property Organization.
- Eloundou, T., Manning, S., Mishkin, P. and Rock, D. (2023). GPTs are GPTs: An early look at the labor market impact potential of large language models. arXiv. Available at: <https://arxiv.org/abs/2303.10130>.
- Elsenbroich, C. (2023). Agent-based modelling as a method for prediction in complex social systems. *International Journal of Social Research Methodology*, 26(2), 133–142. <https://doi.org/10.1080/13645579.2023.2152007>
- ENTSO-E (European Network of Transmission System Operators for Electricity). (2025). European Resource Adequacy Assessment (ERAA) 2024 edition: Annex 3—Detailed results (ACER approved and amended version, August 2025).
- ENTSO-E (European Network of Transmission System Operators for Electricity). (2025). European Resource Adequacy Assessment (ERAA) 2024 edition: Executive report (ACER approved and amended version, August 2025).
- ENTSO-E. (2024). ERAA 2024 Annex 3: Detailed results (v2).

- ENTSO-E. (2025, August). ERAA 2024 Executive Report (ACER-approved amended version).
- Eubanks, V. (2018). Automating inequality: How high-tech tools profile, police, and punish the poor. New York: St. Martin's Press.
- European Commission (2024). The EU AI Act: A Comprehensive Framework for Artificial Intelligence. European Commission Directorate-General for Communications Networks, Content and Technology (CNECT). Available at: <https://digital-strategy.ec.europa.eu/en/policies/european-approach-artificial-intelligence>
- European Commission: European Education and Culture Executive Agency, *Transforming digital education – Insights from three years of the Digital Education Accelerator*, Publications Office of the European Union, 202. <https://data.europa.eu/doi/10.2797/8936638>
- European Commission. (2025, April 9). AI Continent Action Plan (COM(2025) 165 final) https://commission.europa.eu/topics/eu-competitiveness/ai-continent_en
- European Commission. (2025, October 8). A European strategy for artificial intelligence in science: Paving the way for the Resource for AI Science in Europe. https://research-and-innovation.ec.europa.eu/document/c1afd7d0-ff65-4f84-be48-b0e0949596c5_en
- European Commission. (2025, October 8). Apply AI Strategy (COM(2025) 723 final). <https://digital-strategy.ec.europa.eu/en/policies/apply-ai>
- Farmer, J. D. (2025). Quantitative agent-based models: A promising alternative for macroeconomics. Oxford Review of Economic Policy. <https://doi.org/10.1093/oxrep/graf027>
- Federal Reserve Bank of St. Louis. FRED: Economic data. <https://fred.stlouisfed.org/>
- Feenberg, A. (1991). Critical theory of technology. Oxford: Oxford University Press.
- Fernández-Mateo, I. (2025). Generative AI and hiring Journal of Organization Design, 13,4. <https://doi.org/10.1007/s41469-025-00197-1>
- Floridi, L. (2013). The philosophy of information. Oxford University Press.
- Freidson, E. (2001). Professionalism: The third logic. University of Chicago Press.
- Frey, C. B., & Osborne, M. A. (2017). The future of employment: How susceptible are jobs to computerisation? Technological Forecasting and Social Change, 114, 254–280.
- Folke, C., Carpenter, S.R., Walker, B., Scheffer, M., Chapin, T., and Rockström, J. (2010). Resilience thinking: Integrating resilience, adaptability and transformability. Ecology and Society, 15(4), p.20. Available at: <https://www.ecologyandsociety.org/vol15/iss4/art20/>.
- Floridi, L. and Cows, J. (2019). A unified framework of five principles for AI in society. Harvard Data Science Review, 1(1). Available at: <https://hdsr.mitpress.mit.edu/pub/ph1o75no/release/>.
- Freedom House. (2024). Freedom on the Net 2024: The struggle for trust online. <https://freedomhouse.org/report/freedom-net/2024/struggle-trust-online>
- Freedom House. (2025). Freedom in the World 2025: The uphill battle to safeguard rights. <https://freedomhouse.org/report/freedom-world/2025/uphill-battle-to-safeguard-rights>
- Friis, A., & Riley, J. (2025). Performance or principle? Resistance to AI in the U.S. labor market. https://www.hbs.edu/ris/Publication%20Files/26-017_5411629b-87ca-463d-b7ab-534032419c6b.pdf
- Georgieva, K. (2025, October 8). Opportunity in a time of change (IMF Annual Meeting 2025 curtain-raiser). <https://www.imf.org/en/News/Articles/2025/10/08/sp100825-annual-meetings-2025-curtain-raiser>
- Gerlich, M. (2025). AI tools in society: Impacts on cognitive offloading and the future of critical thinking. Societies, 15(1), 6. <https://www.mdpi.com/2075-4698/15/1/6>
- Glăveanu, V. P. (2013). Rewriting the language of creativity: The five A's framework. Review of General Psychology, 17(1), 69–81. <https://doi.org/10.1037/a0029528>

Glăveanu, V. P. (2020). A sociocultural theory of creativity: Bridging the social, the material, and the psychological. *Review of General Psychology*, 24(4), 335–354. <https://psycnet.apa.org/record/2020-85633-003>

Gonsalves, R. (2024). Generative AI's impact on critical thinking: Revisiting Bloom's taxonomy. <https://journals.sagepub.com/doi/10.1177/02734753241305980>

Gorupec, N., Brehmer, N., Tiberius, V., & Kraus, S. (2022). Tackling uncertain future scenarios with real options: A review and research framework. *Irish Journal of Management*, 41(1), 69–88. <https://doi.org/10.2478/ijm-2022-0003>

Guilford, J. P. (1950). Creativity. *American Psychologist*, 5(9), 444–454. <https://doi.org/10.1037/h0063487>

Hasson, F., Keeney, S., & McKenna, H. (2025). Revisiting the Delphi technique—Research thinking and practice: A discussion paper. *International Journal of Nursing Studies*, 168, 105119. <https://doi.org/10.1016/j.ijnurstu.2025.105119>

HEPI (Higher Education Policy Institute). (2025). AI and the future of universities. <https://www.hepi.ac.uk/wp-content/uploads/2025/10/AI-and-the-Future-of-Universities.pdf>

Hernández-Ramos, J., Araya, R. (2025). Do school activities foster creative thinking? An analysis of PISA results. *Education Sciences*, 15(2), 133. <https://doi.org/10.3390/educsci15020133>

Hinrich Foundation & IMD World Competitiveness Center. (2025). Sustainable trade index 2025: Navigating sustainable trade. From benchmarking to trade-offs. Hinrich Foundation and IMD. <https://www.imd.org/centers/wcc/world-competitiveness-center/rankings/sustainable-trade-index/>

Holling, C.S. (1973). Resilience and stability of ecological systems. *Annual Review of Ecology and Systematics*, 4, pp.1–23. Available at: <https://doi.org/10.1146/annurev.es.04.110173.000245>.

HoloniQ. (2025). Education in 2030: Five scenarios for the future of learning and talent. <https://www.holoniq.com/2030>

Hosseini, S.M., Lichtinger, G. (2025) Generative AI as Seniority-Biased Technological Change: Evidence from U.S. Resume and Job Posting Data. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5425555

IMD (International Institute for Management Development). (2025). The future of supply chain 2025. <https://www.imd.org/ibyimd/white-papers/the-future-of-supply-chain/>

Tan, J., Jackson, B., Berjon, R. & Coyle, D. (2025). Airbus for AI: A global strategy for public value creation. Bennett School for Public Policy, University of Cambridge. <https://doi.org/10.17863/CAM.121150>

Intergovernmental Panel on Climate Change. (2023). Climate change 2023: Synthesis report (H. Lee & J. Romero, Eds.). IPCC. <https://www.ipcc.ch/report/ar6/syr/>

International AI Safety Community. (2025). International AI Safety 2025 update. https://internationalaisafetyreport.org/sites/default/files/2025-10/international_ai_safety_report_2025_english.pdf

International Energy Agency. (2023). Electricity grids and secure energy transitions: Enhancing the foundations of resilient, sustainable and affordable power systems. <https://iea.blob.core.windows.net/assets/ea2ff609-8180-4312-8de9-494bcf21696d/ElectricityGridsandSecureEnergyTransitions.pdf>

International Energy Agency. (2024). Renewables 2024: Analysis and forecast to 2030. <https://www.iea.org/reports/renewables-2024>

International Labour Organization. (2025). World employment and social outlook: Trends 2025 https://www.ilo.org/sites/default/files/2025-01/WESO25_Trends_EN_WEB5.pdf

International Monetary Fund (IMF). (2024). Gen-AI: Artificial Intelligence and the Future of Work. Washington, DC: IMF Publications: <https://www.imf.org/en/Publications/Staff-Discussion-Notes/Issues/2024/01/14/Gen-AI-Artificial-Intelligence-and-the-Future-of-Work-542379>.

International Monetary Fund. (2024, October). World Economic Outlook: Policy pivot, rising threats. <https://www.imf.org/en/Publications/WEO/Issues/2024/10/22/world-economic-outlook-october-2024>

International Monetary Fund. (2025, April). World Economic Outlook: A critical juncture amid policy shifts. <https://www.imf.org/en/Publications/WEO/Issues/2025/04/22/world-economic-outlook-april-2025>

International Monetary Fund. (2025, January). World Economic Outlook Update: Global growth—Divergent and uncertain. International Monetary Fund. <https://www.imf.org/en/Publications/WEO/Issues/2025/01/17/world-economic-outlook-update-january-2025>

International Monetary Fund. (2025, July). World Economic Outlook Update: Global economy-Tenuous resilience amid persistent uncertainty. International Monetary Fund. <https://www.imf.org/en/Publications/WEO/Issues/2025/07/29/world-economic-outlook-update-july-2025>

International Monetary Fund. (2025, October). World Economic Outlook: Global economy in flux, prospects remain dim. <https://www.imf.org/en/Publications/WEO/Issues/2025/10/14/world-economic-outlook-october-2025>

International Monetary Fund. International Monetary Fund. IMF data. <https://www.imf.org/en/Data>

Jeffrey, B., Craft, A. (2004). Teaching creatively and teaching for creativity: Distinctions and relationships. *Educational Studies*, 30(1), 77–87. <https://doi.org/10.1080/0305569032000159750>

Jessop, B. (2002). Liberalism, neo-liberalism, and urban governance: A state-theoretical perspective. *Antipode*, 34(3), pp.452-472.

Kahneman, D. (2021). 'Clearly AI is going to win. How people are going to adjust is a fascinating problem.' *The Guardian*. Available at: <https://www.theguardian.com/books/2021/may/16/daniel-kahneman-clearly-ai-is-going-to-win-how-people-are-going-to-adjust-is-a-fascinating-problem-thinking-fast-and-slow>

Kamar, E. (2016). Directions in hybrid intelligence: Complementing AI systems with human intelligence. In *Proceedings of the Twenty-fifth International Joint Conference on Artificial Intelligence (IJCAI)* (pp. 4070–4073).

Kaplan, J., McCandlish, S., Henighan, T., et al. (2020). Scaling laws for neural language models. *arXiv*. <https://arxiv.org/abs/2001.08361>

Kaufman, J. C., Sternberg, R. J. (2019). *The Cambridge Handbook of Creativity* (2nd ed.). Cambridge University Press. <https://doi.org/10.1017/9781316979839>

Kharkhurin, A. V. (2014). Creativity. 4in1: Four-criterion construct of creativity. *Creativity Research Journal*, 26(3), 338–352. DOI: [10.1080/10400419.2014.929424](https://doi.org/10.1080/10400419.2014.929424)

Kim, K. H. (2006). Can we trust creativity tests? A review of the Torrance Tests of Creative Thinking (TTCT). *Creativity Research Journal*, 18(1), 3–14. https://doi.org/10.1207/s15326934crj1801_2

Kosmyna, N. (2025). Your brain on ChatGPT: Accumulation of cognitive debt when using an AI assistant for essay writing task. MIT Media Lab. DOI: [10.48550/arXiv.2506.08872](https://doi.org/10.48550/arXiv.2506.08872)

Kurniawan, J. H., Apergi, M., Eicke, L., Goldthau, A., Lazurko, A., Nordemann, E., Schuch, E., Sharma, A., Siddhantakar, N., Veit, K., & Weko, S. (2022). Towards participatory cross-impact balance analysis: Leveraging morphological analysis for data collection in energy transition scenario workshops. *Energy Research & Social Science*, 93, 102815. <https://doi.org/10.1016/j.erss.2022.102815>

Lane, M. (2024). Who will be the workers most affected by AI? A closer look at the impact of AI on women, low-skilled workers and other groups (OECD Artificial Intelligence Papers No. 26). Organization for Economic Co-operation and Development. https://www.oecd.org/content/dam/oecd/en/publications/reports/2024/10/who-will-be-the-workers-most-affected-by-ai_fb7fcccd/14dc6f89-en.pdf

- LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436–444.
<https://doi.org/10.1038/nature14539>
- LinkedIn. (2023). Future of Work Report: AI in Work.
<https://economicgraph.linkedin.com/content/dam/me/economicgraph/en-us/PDF/future-of-work-report-ai-november-2023.pdf>
- Lucas, B. (2022). A field guide to assessing creative thinking in schools. DOI: [10.13140/RG.2.2.24010.03529](https://doi.org/10.13140/RG.2.2.24010.03529)
- Lucas, B. (2025). The development of a conceptual framework for embedding creativity in schools. *Thinking Skills and Creativity*, 58, 101874. <https://doi.org/10.1016/j.tsc.2025.101874>
- Lucas, B., Claxton, G., Spencer, E. (2013). Progression in student creativity in school: First steps towards new forms of formative assessments. OECD.
https://www.oecd.org/content/dam/oecd/en/publications/reports/2013/01/progression-in-student-creativity-in-school_g17a2218/5k4dp59msdwk-en.pdf
- Luckin, R. (2018). Machine learning and human intelligence. UCL IOE Press.
- Makridakis, S., Spiliotis, E., & Assimakopoulos, V. (2018). Statistical and machine learning forecasting methods: Concerns and ways forward. *PLOS ONE*, 13(3), e0194889.
<https://doi.org/10.1371/journal.pone.0194889>
- Manca, F. and D. Eslava (2025), "Digital and AI skills in health occupations: What do we know about new demand?", *OECD Artificial Intelligence Papers*, No. 36, OECD Publishing, Paris, <https://doi.org/10.1787/5fbd42ab-en>.
- Manning, J., & Horton, J. (2025). General Social Agents. arXiv:2508.17047v1
- Martello, M. V., Whittle, A. J., Oddo, P. C., & de Neufville, R. (2024). Real options analysis for valuation of climate adaptation pathways with application to transit infrastructure. *Risk Analysis*, 44(5), 1046–1066.
<https://doi.org/10.1111/risa.14218>
- McKee, N.J., Kozleski, E., Lemons, C.J., Hau, I.C. (2025) AI + learning differences: Designing a future with no boundaries. Stanford Accelerator for Learning.
https://acceleratelearning.stanford.edu/app/uploads/2025/07/AI-Learning-Differences-Designing-a-Future-with-No-Boundaries_Final.pdf
- McKinsey & Company. (2025). Extracting value from AI in banking: Rewiring the enterprise.
<https://www.mckinsey.com/industries/financial-services/our-insights/extracting-value-from-ai-in-banking-rewiring-the-enterprise>
- McKinsey and Company. (2023). Generative AI and the Future of Work in America. McKinsey Global Institute. <https://www.mckinsey.com/mgi/our-research/generative-ai-and-the-future-of-work-in-america#/>
- McKinsey Global Institute. (2023). The Economic Potential of Generative AI: Pioneering Transformation across Industries. <https://www.mckinsey.com/capabilities/tech-and-ai/our-insights/the-economic-potential-of-generative-ai-the-next-productivity-frontier#/>
- Microsoft (2025) AI in Education: A Special Report. <https://cdn-dynmedia-1.microsoft.com/is/content/microsoftcorp/microsoft/bade/documents/products-and-services/en-us/education/2025-Microsoft-AI-in-Education-Report.pdf>
- Milanez, A., A. Lemmens and C. Ruggiu (2025), "Algorithmic management in the workplace: New evidence from an OECD employer survey", *OECD Artificial Intelligence Papers*, No. 31, OECD Publishing, Paris, <https://doi.org/10.1787/287c13c4-en>.
- Ministry of Education of the People's Republic of China. (2025). 国家中小学智慧教育平台与人工智能融合应用指南（试行） [National primary and secondary smart education platform and AI integration application guide (trial)]. Ministry of Education of the PRC.

- Misch F., Park, B., Pizzinelli, C., SherG. "AI and Productivity in Europe", IMF Working Papers 2025, 067 (2025). <https://doi.org/10.5089/9798229006057.001>
- MIT Technology Review Insights. (2023). The great acceleration: CIO perspectives on generative AI. <https://www.technologyreview.com/2023/07/18/1076423/the-great-acceleration-cio-perspectives-on-generative-ai/>
- Mollick, E. (2024). Co-intelligence: Living and working with AI. Portfolio/Penguin.
- NATO Allied Command Transformation. (2023). NATO Wargaming Handbook. Norfolk, VA. <https://www.act.nato.int/>
- Newman, N., Ross Arguedas, A., Robertson, C. T., Nielsen, R. K., & Fletcher, R. (2025). Digital News Report 2025. Reuters Institute for the Study of Journalism. <https://reutersinstitute.politics.ox.ac.uk/digital-news-report/2025>
- Ng, A. (2018). Machine Learning Yearning. DeepLearning.AI. Available at: <https://info.deeplearning.ai/machine-learning-yearning-book>.
- NITI Aayog. (2021). National Strategy for Artificial Intelligence: #AIforAll. New Delhi: Government of India. <https://www.niti.gov.in/sites/default/files/2023-03/National-Strategy-for-Artificial-Intelligence.pdf>.
- Niu, W., Sternberg, R. J. (2006). The philosophical roots of Western and Eastern conceptions of creativity. *Journal of Theoretical and Philosophical Psychology*, 26(1-2), 18–38. <https://doi.org/10.1037/h0091265>
- North American Electric Reliability Corporation. (2024). 2024 long-term reliability assessment (Updated July 15, 2025). North American Electric Reliability Corporation. https://www.nerc.com/pa/RAPA/ra/Reliability%20Assessments%20DL/NERC_Long%20Term%20Reliability%20Assessment_2024.pdf
- Noy, S., & Zhang, W. (2023). Experimental Evidence on the Productivity Effects of Generative AI. *Science* 381.6654 <https://www.science.org/doi/10.1126/science.adh2586>
- OECD (2025) Recommendation of the Council on Artificial Intelligence, available at <https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0449>
- Organization for Economic Co-operation and Development (OECD). (2024). Framework for Anticipatory Governance of Emerging Technologies, *OECD Science, Technology and Industry Policy Papers*, No. 165, OECD Publishing, Paris, <https://doi.org/10.1787/0248ead5-en>.
- Office for National Statistics. Datasets. <https://www.ons.gov.uk/datasets>
- Organization for Economic Co-operation and Development (OECD) Data. <https://data.oecd.org/>
- Organization for Economic Co-operation and Development (OECD) (2019). Future of Education and Skills 2030: OECD Learning Compass 2030. OECD Publishing. <https://www.oecd.org/en/data/tools/oecd-learning-compass-2030.html>
- Organization for Economic Co-operation and Development (OECD) (2023). PISA 2022 Results (Volume I): The State of Learning and Equity in Education, PISA, OECD Publishing, Paris, <https://doi.org/10.1787/53f23881-en>
- Organization for Economic Co-operation and Development (OECD) (2018). The Future of Education and Skills: Education 2030, *OECD Education Policy Perspectives*, No. 98, OECD Publishing, Paris, <https://doi.org/10.1787/54ac7020-en>.
- Organization for Economic Co-operation and Development, & United Nations Educational, Scientific and Cultural Organization. (2024). G7 Toolkit for Artificial Intelligence in the Public Sector, OECD Publishing, Paris, <https://doi.org/10.1787/421c1244-en>.
- Organization for Economic Co-operation and Development. (2023). OECD Employment Outlook 2023: Artificial Intelligence and the Labour Market, OECD Publishing, Paris, <https://doi.org/10.1787/08785bba-en>.

Organization for Economic Co-operation and Development. (2023). OECD Science, Technology and Innovation Outlook 2023: Enabling Transitions in Times of Disruption, OECD Publishing, Paris, <https://doi.org/10.1787/0b55736e-en>.

Organization for Economic Co-operation and Development. (2024). Education Policy Outlook 2024: Reshaping Teaching into a Thriving Profession from ABCs to AI, OECD Publishing, Paris, <https://doi.org/10.1787/dd5140e4-en>.

Organization for Economic Co-operation and Development. (2024). OECD Employment Outlook 2024: The Net-Zero Transition and the Labour Market, OECD Publishing, Paris, <https://doi.org/10.1787/ac8b3538-en>.

Organization for Economic Co-operation and Development. (2024). PISA 2022 Results (Volume V): Learning Strategies and Attitudes for Life, PISA, OECD Publishing, Paris, <https://doi.org/10.1787/c2e44201-en>.

Organization for Economic Co-operation and Development. (2025). AI openness: A primer for policymakers, OECD Artificial Intelligence Papers, No. 44, OECD Publishing, Paris, <https://doi.org/10.1787/02f73362-en>.

Organization for Economic Co-operation and Development. (2025). Education at a Glance 2025: OECD Indicators, OECD Publishing, Paris, <https://doi.org/10.1787/1c0d9c79-en>.

Organization for Economic Co-operation and Development. (2025). Empowering the Workforce in the Context of a Skills-First Approach, OECD Skills Studies, OECD Publishing, Paris, <https://doi.org/10.1787/345b6528-en>.

Organization for Economic Co-operation and Development. (2025). OECD (2025), Introducing the OECD AI Capability Indicators, OECD Publishing, Paris, <https://doi.org/10.1787/be745f04-en>.

Organization for Economic Co-operation and Development. (2025). OECD Employment Outlook 2025: Can We Get Through the Demographic Crunch?, OECD Publishing, Paris, <https://doi.org/10.1787/194a947b-en>.

Organization for Economic Co-operation and Development. (2025). Trends Shaping Education 2025, OECD Publishing, Paris, <https://doi.org/10.1787/ee6587fd-en>.

Organization for Economic Co-operation and Development. (2025). The Global Forum on Productivity at 10: Past and Future Perspectives on Reviving Productivity Growth, OECD Publishing, Paris, <https://doi.org/10.1787/ca7295d9-en>.

Ostrom E. (2009) A general framework for analyzing sustainability of social-ecological systems. *Science*. 2009 Jul 24;325(5939):419-22. DOI: 10.1126/science.1172133. PMID: 19628857.

Oxford Insights. (2024). Government AI Readiness Index 2024. Oxford: Oxford Insights. <https://oxfordinsights.com/ai-readiness/ai-readiness-index/>.

Oxford University Press. (2025). Teaching the AI-native generation. Oxford University Press.

Phaal, R. (2024). Strategic planning for roadmapping systems (Working paper). University of Cambridge, Institute for Manufacturing. <https://engage.ifm.eng.cam.ac.uk/project/an-introductory-guide-to-strategic-and-technology-roadmaps-2/>

Wharton Human-AI Research and GBK Collective. https://ai.wharton.upenn.edu/wp-content/uploads/2025/10/2025-Wharton-GBK-AI-Adoption-Report_Full-Report.pdf

Rahwan, I. (2018). Society-in-the-loop: Programming the algorithmic social contract. *Ethics of Information Technology*, 20(5), pp.5-14. Available at: <https://doi.org/10.1007/s10676-017-9430-8>.

Reiter-Palmon, R., Forthmann, B., Barbot, B. (2019). Scoring divergent thinking tests: A review and systematic framework. *Psychology of Aesthetics, Creativity, and the Arts*, 13(2), 144–152. <https://doi.org/10.1037/aca0000227>

Restrepo, P (2025) We Won't be Missed: Work and Growth in the AGI World <https://www.nber.org/system/files/chapters/c15315/c15315.pdf>

- Riedl, M., & Weidmann, N. (2025). Quantifying Human-AI synergy. https://osf.io/preprints/psyarxiv/vbkmt_v1
- Risko, E. F., & Gilbert, S. J. (2016). Cognitive offloading. *Trends in Cognitive Sciences*, 20(9), 676–688. <https://doi.org/10.1016/j.tics.2016.07.002>
- Ritchey, T. (2022). General morphological analysis: An overview. *Academia Letters*, Article 4620. <https://doi.org/10.20935/AL4620>
- Robinson, K. (2011). *Out of our minds: Learning to be creative* (2nd ed.). Capstone.
- Rodrik, D. (2011). The future of economic convergence. NBER Working Paper No. 17400. DOI 10.3386/w17400
- Runco, M. A., Jaeger, G. J. (2012). The standard definition of creativity. *Creativity Research Journal*, 24(1), 92–96. <https://doi.org/10.1080/10400419.2012.650092>
- Russell, S., & Norvig, P. (2021). *Artificial intelligence: A modern approach* (4th ed.). Pearson.
- Saar, A.B. and Busuioc, M. (2023). Human-AI interactions in public sector decision making: Automation bias and selective adherence to algorithmic advice. *Journal of Public Administration Research and Theory*, 33(1), pp.153-169. DOI: [10.30636/jbpa.71.261](https://doi.org/10.30636/jbpa.71.261)
- Saeed, B. A., Ramdane, T. M. (2022). The effect of implementation of a creative thinking model on the development of creative thinking skills in high school students: A systematic review. *Review of Education*, 10(4), e3379. <https://doi.org/10.1002/rev3.3379>
- Schifano, J., & Niederberger, M. (2025). How Delphi studies in the health sciences find consensus: A scoping review. *Systematic Reviews*, 14, 14. <https://doi.org/10.1186/s13643-024-02738-3>
- Schoemaker, P. J. H. (1995). Scenario planning: A tool for strategic thinking. *Sloan Management Review*, 36(2), 25–40. https://www.researchgate.net/publication/220042263_Scenario_Planning_A_Tool_for_Strategic_Thinking
- Schmidt, J., G. Pilgrim and A. Mourougane (2024), Measuring the demand for AI skills in the United Kingdom, OECD Artificial Intelligence Papers, No. 25, OECD Publishing, Paris, <https://doi.org/10.1787/1d6474ef-en>.
- Schön, D. A. (1983). *The reflective practitioner: How professionals think in action*. Basic Books.
- Schünemann, C., Johanning, S., Reger, E., Herold, H., & Bruckner, T. (2024). Complex system policy modelling approaches for policy advice—Comparing systems thinking, system dynamics and agent-based modelling. *Political Research Exchange*, 6, 2387438. <https://doi.org/10.1080/2474736X.2024.2387438>
- Schwartz, P. (1996). *The art of the long view: Planning for the future in an uncertain world*. Currency/Doubleday.
- Sevilla, J., Heim, L., Hobhahn, M., et al. (2022). Compute trends across three eras of machine learning. *arXiv* [10.48550/arXiv.2202.05924](https://arxiv.org/abs/2202.05924)
- Simpson, E. (2025) Empowerment or Dependency? A Research Agenda for Human Agency in the Age of AI. DOI: <http://dx.doi.org/10.2139/ssrn.5401467>.
- Simpson, E (2025) The AI Matrix: Empowerment or Dependency? A Conceptual Framework. DOI: <http://dx.doi.org/10.2139/ssrn.5228571>
- Simpson, Ewan, The AI Passport: Towards A New Conceptual Framework For Global Skills Certification. DOI: <http://dx.doi.org/10.2139/ssrn.5343668>
- Simpson, E. (2025) Creativity and the AI Era: Why PISA 2022 Matters and What Comes Next: A Research Agenda. DOI: <http://dx.doi.org/10.2139/ssrn.5416896>.

Simpson, E. (2025) A Generation at Risk? Creativity, PISA 2022, and the Demands of an AI Economy and Society in 2025. <https://ssrn.com/abstract=5552659>.

Smart Nation Singapore (2023). Smart Nation 2 Report: A Thriving Digital Future for All. <https://file.go.gov.sg/smartnation2-report.pdf>

Spaltini, M., Acerbi, F., De Carolis, A., Terzi, S., & Taisch, M. (2024). Toward a technology roadmapping methodology to enhance sustainable and digital transition in manufacturing. *Production & Manufacturing Research*, 12(1), 614–649. <https://doi.org/10.1080/21693277.2023.2298572>

Sparrow B, Liu J, Wegner DM. (2011)Google effects on memory: cognitive consequences of having information at our fingertips. *Science*. 2011 Aug 5;333(6043):776-8. doi: 10.1126/science.120774

Squicciarini, M. and H. Nachtigall (2021), "Demand for AI skills in jobs: Evidence from online job postings", OECD Science, Technology and Industry Working Papers, No. 2021/03, OECD Publishing, Paris, <https://doi.org/10.1787/3ed32d94-en>.

Stanford Institute for Human-Centered Artificial Intelligence (HAI). (2023). AI Index Report 2023. Stanford, CA: Stanford University. Available at: <https://aiindex.stanford.edu/report/2023/>.

Stanford Institute for Human-Centered Artificial Intelligence (HAI). (2024). AI Index Report 2024. Stanford, CA: Stanford University. Available at: <https://aiindex.stanford.edu/report/2024/>.

Stanford Institute for Human-Centered Artificial Intelligence (HAI) (2025). 2025 AI Index Report. Stanford Human-Centered Artificial Intelligence. <https://hai.stanford.edu/ai-index/2025-ai-index-report>

Susskind, R., & Susskind, D. (2015). *The future of the professions: How technology will transform the work of human experts*. Oxford University Press.

Tetlock, P. E., & Gardner, D. (2015). *Superforecasting: The art and science of prediction*. Crown.

The Alan Turing Institute. (2024). Understanding the impacts of generative AI use on children. The Alan Turing Institute. <https://www.turing.ac.uk/research/research-projects/understanding-impacts-generative-ai-use-children>.

Torrance, E. P. (1966). *Torrance Tests of Creative Thinking*. Personnel Press.

Tortoise Media. (2023). *The Global AI Index 2023*. London: Tortoise Media. Available at: <https://www.tortoisemedia.com/data/global-ai/>.

Trilling, B., Fadel, C. (2009). *21st century skills: Learning for life in our times*. Jossey-Bass.

U.S. Department of Education. (2023). *Artificial intelligence and the future of teaching and learning: Insights and recommendations*. <https://www.ed.gov/sites/ed/files/documents/ai-report/ai-report.pdf>

UK Government Office for Science. (2024). *The Futures Toolkit (2nd ed.)*. London: HM Government. <https://www.gov.uk/government/publications/futures-toolkit-for-policy-makers-and-analysts/the-futures-toolkit-html>

UK Government Office for Science. (2025). *A brief guide to futures thinking and foresight*. London: HM Government. https://assets.publishing.service.gov.uk/media/68ed18e382670806f9d5dff6/A_Brief_Guide_to_Futures_Thinking_and_Foresight_2025.pdf

UK Ministry of Defence (Development, Concepts and Doctrine Centre). (2023). *Influence Wargaming Handbook*. London: Ministry of Defence. https://assets.publishing.service.gov.uk/media/6494481b9e7a8b00139329d8/Influence_Wargaming_Handbook_web.pdf

UNESCO. (2021). *Reimagining our futures together: A new social contract for education*. UNESCO. <https://unesdoc.unesco.org/ark:/48223/pf0000379707>

United Nations. (2024). *Governing AI for humanity: Final report*. United Nations. United Nations. https://www.un.org/sites/un2.un.org/files/governing_ai_for_humanity_final_report_en.pdf

United Nations (n.d.) Datasets. <https://data.un.org/>

van der Heijden, K. (1996). *Scenarios: The art of strategic conversation*. Wiley.

van der Voorn, T., Quist, J., Svenfelt, Å., Kok, K., Rousval, A., Bhadwal, S., Carmichael, J., Hatzilakou, D. (2023). Advancing participatory backcasting for climate change adaptation planning using ten cases from three continents. *Climate Risk Management*, 41, 100559. <https://doi.org/10.1016/j.crm.2023.100559>

Vaswani, A., Shazeer, N., Parmar, N., et al. (2017). Attention is all you need. *Advances in Neural Information Processing Systems*, 30. <https://dl.acm.org/doi/pdf/10.5555/3295222.3295349>

Wack, P. (1985). Scenarios: Uncharted waters ahead. *Harvard Business Review*, 63(5), 73–89. <https://hbr.org/1985/09/scenarios-uncharted-waters-ahead>

Weimer-Jehle, W. (2006). A system-theoretical approach to cross-impact analysis. *Technological Forecasting & Social Change*, 73(4), 334–361. <https://doi.org/10.1016/j.techfore.2005.06.005>

Weimer-Jehle, W. (2023). *Cross-Impact Balances (CIB) for scenario analysis: Fundamentals and implementation*. Cham: Springer. <https://doi.org/10.1007/978-3-031-27230-1>

Wilson, H. J., & Daugherty, P. R. (2018). Collaborative intelligence: Humans and AI are joining forces. *Harvard Business Review*, 96(4), 114–123. <https://hbr.org/2018/07/collaborative-intelligence-humans-and-ai-are-joining-forces>

World Bank. (2022). *Equitable human capital development in Kazakhstan: Technical and vocational skills development, adult learning and labor market programs*. World Bank. <https://documents1.worldbank.org/curated/en/099083002082333944/pdf/P17487900b2b230ae09b920ba9aec6d07c9.pdf>

World Bank. 2025. *Building Trustworthy Artificial Intelligence: Frameworks, Applications, and Self-Assessment for Readiness*. World Bank. <http://hdl.handle.net/10986/43656>

World Bank Group. 2025. *Devising a Strategic Approach to Artificial Intelligence: A Handbook for Policy Makers*. World Bank. <http://hdl.handle.net/10986/43347>

World Bank. (2025). *Global economic prospects: June 2025*. World Bank. <https://openknowledge.worldbank.org/server/api/core/bitstreams/0e685254-776a-40cf-b0ac-f329dd182e9b/content>

World Bank (2025). *World Bank Open Data*. <https://data.worldbank.org/>

World Economic Forum. (2023). *Future of Jobs Report 2023*. Geneva: WEF Publications. Available at: <https://www.weforum.org/reports/the-future-of-jobs-report-2023/>.

World Economic Forum. (2025). *Future of jobs report 2025*. World Economic Forum. <https://www.weforum.org/publications/the-future-of-jobs-report-2025/>

World Economic Forum (2025). *10 emerging technology solutions for planetary health*. <https://www.weforum.org/publications/10-emerging-technology-solutions-for-planetary-health/>

World Economic Forum (2025). *Quantum technologies: Key strategies and opportunities for financial services leaders*. https://reports.weforum.org/docs/WEF_Quantum_Technologies_Key_Strategies_and_Opportunities_for_Financial_Services_Leaders_2025.pdf

Wübbelmann, T., & Kabisch, N. (2025). Backcasting—A scenario approach in urban climate adaptation planning. *npj Urban Sustainability*, 5, 69. <https://doi.org/10.1038/s42949-025-00260-6>. DOI: [10.1002/berj.3135](https://doi.org/10.1002/berj.3135)

Zuboff, S. (2019). *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*. New York: Public Affairs.

Technical Appendix 1 : The “Historic Canon”

The historic canon is a concise map of the big ideas humanity has developed about people, power, knowledge and limits, organized so they can guide how we build and govern modern technologies. Instead of treating artificial intelligence as something without precedent, the canon gathers the most relevant strands of global thought and shows how they speak to today’s decisions.

From philosophy and political theory (Plato, Aristotle, Kant, Rawls, Arendt) we take questions about autonomy, justice and the ‘good life’. Religious and ethical traditions (Abrahamic, Eastern and Indigenous) contribute duties of stewardship, restraint and compassion. Psychology and cognitive science (Piaget, Vygotsky, Kahneman, Simon) explain bounded reasoning and learning, reminding us that both humans and machines operate with limits. Education thinkers (Socrates, Confucius, Dewey, Freire, hooks, Montessori) keep purpose, dialogue and formation at the center, not just efficiency.

Political economy and management (Smith, Marx, Polanyi, Schumpeter, Drucker, Piketty, Mazzucato) explain how technology reshapes work, ownership and inequality. Media and communication theory (Lasswell, Lippmann, McLuhan, Hall, Zuboff) show how information systems manufacture consent, create filter bubbles and centralize power. Science and technology studies, anthropology and linguistics (Kuhn, Latour, Haraway, Geertz, Said, Saussure, Chomsky) reveal how knowledge and artifacts are made within culture. Environmental thought (Leopold, Carson, Lovelock, Naess) brings planetary limits and ecological justice to the fore. Decolonial and Global South voices (Fanon, Ngũgĩ, Quijano, Dussel, Santos, Mbembe, Shiva) expose extraction, epistemic inequality and the need for cognitive justice. Technical AI voices (Turing, Wiener, McCarthy, Minsky, Pearl, Hinton, LeCun, Bengio, Russell, Gebru, Li) provide the field’s own breakthroughs and cautions.

Operationally, the canon does three things. First, it makes first principles explicit and versioned, so judgments are traceable. Second, it translates principles into practical guardrails (what we count, what clears a threshold, what triggers reversion). Third, it widens the lens beyond short-term metrics to equity, legitimacy and ecological fit. In short, the canon is the backbone of the Grey Swan framework’s Tier 0 rulebook: a shared, teachable foundation that keeps evidence, ethics and agency aligned as we assess futures.

Canon References

I. Fundamentals of Human Thought and Meaning

Aquinas, T. (2006). *Selected philosophical writings* (T. McDermott, Ed. & Trans.). Oxford University Press. (Original works published 13th century)

Aristotle. (2009). *Nicomachean ethics* (W. D. Ross, Trans.). Oxford University Press. (Original work published ca. 350 BCE)

Augustine. (1991). *Confessions* (H. Chadwick, Trans.). Oxford University Press. (Original work published ca. 397 CE)

Bhagavad Gītā. (2000). *The Bhagavad Gita* (E. Easwaran, Trans.). Nilgiri Press. (Original work composed ca. 2nd c. BCE–2nd c. CE)

Cicero. (2016). *On the ideal orator* (J. M. May & J. Wisse, Trans.). Oxford University Press. (Original work published ca. 55 BCE)

Confucius. (1998). *The analects* (R. Dawson, Trans.). Oxford University Press. (Original work compiled ca. 5th–3rd c. BCE)

Laozi. (2009). *Daodejing* (P. J. Ivanhoe, Trans.). Hackett. (Original work composed ca. 4th–3rd c. BCE)

Plato. (1997). *Phaedrus* (A. Nehamas & P. Woodruff, Trans.). In J. M. Cooper (Ed.), *Plato: Complete works* (pp. 506–556). Hackett. (Original work published ca. 370 BCE)

Plato. (1997). *Republic* (G. M. A. Grube, Trans., rev. C. D. C. Reeve). Hackett. (Original work published ca. 375 BCE)

Quintilian. (2001). *The orator's education* (D. A. Russell, Trans.; 5 vols.). Harvard University Press.

Zhuangzi. (2009). *The essential Zhuangzi* (S. Ziporyn, Trans.). Hackett. (Original work composed ca. 4th–3rd c. BCE)

II. Knowledge, Mind, and Learning

Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, M. (2021). On the dangers of stochastic parrots: Can language models be too big? *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FAccT'21)*, 610–623.

Bernstein, B. (1975). *Class, codes and control: Towards a theory of educational transmissions* (Vol. 3). Routledge.

Biesta, G. (2010). *Good education in an age of measurement: Ethics, politics, democracy*. Routledge.

- Brooks, R. A. (1991). Intelligence without representation. *Artificial Intelligence*, 47(1–3), 139–159.
- Dewey, J. (1916). *Democracy and education*. Macmillan.
- Feyerabend, P. (1975). *Against method: Outline of an anarchistic theory of knowledge*. London: Verso.
- Froebel, F. (2001). *The education of man* (W. N. Hailmann, Trans.). Dover. (Original work published 1826)
- Harding, S. (1991). *Whose science? Whose knowledge? Thinking from women's lives*. Ithaca, NY: Cornell University Press.
- hooks, b. (1994). *Teaching to transgress: Education as the practice of freedom*. Routledge.
- Illich, I. (1971). *Deschooling society*. Harper & Row.
- Kant, I. (1996). *Practical philosophy* (M. J. Gregor, Ed. & Trans.). Cambridge University Press. (Original works published 1785–1797)
- Khan, S. (2012). *The one world schoolhouse: Education reimagined*. Twelve.
- Kuhn, T. S. (1962). *The structure of scientific revolutions*. Chicago: University of Chicago Press.
- Lakatos, I. (1978). *The methodology of scientific research programmes*. Cambridge: Cambridge University Press.
- Li, F.-F., Deng, J., & Fei-Fei, L. (2009). ImageNet: A large-scale hierarchical image database. *IEEE Conference on Computer Vision and Pattern Recognition*, 248–255.
- Longino, H. E. (1990). *Science as social knowledge: Values and objectivity in scientific inquiry*. Princeton, NJ: Princeton University Press.
- McCarthy, J. (1960). Recursive functions of symbolic expressions and their computation by machine, Part I. *Communications of the ACM*, 3(4), 184–195.
- Minsky, M., & Papert, S. (1969). *Perceptrons*. MIT Press.
- Mitra, S. (2012). *Beyond the hole in the wall: Discover the power of self-organized learning*. TED Books.
- Montessori, M. (1964). *The Montessori method*. Schocken. (Original work published 1912)
- Newell, A., & Simon, H. A. (1972). *Human problem solving*. Prentice Hall.
- Pearl, J. (2000). *Causality: Models, reasoning, and inference*. Cambridge University Press.
- Pestalozzi, J. H. (2005). *How Gertrude teaches her children* (L. Holland & F. C. Turner, Trans.). Dover. (Original work published 1801)

- Popper, K. (1959). *The logic of scientific discovery*. London: Routledge.
- Rousseau, J.-J. (1979). *Emile, or on education* (A. Bloom, Trans.). Basic Books. (Original work published 1762)
- Russell, S. (2019). *Human compatible: Artificial intelligence and the problem of control*. Viking.
- Selwyn, N. (2016). *Education and technology: Key issues and debates* (2nd ed.). Bloomsbury.
- Turing, A. M. (1950). Computing machinery and intelligence. *Mind*, 59(236), 433–460.
- Watters, A. (2021). *Teaching machines: The history of personalized learning*. MIT Press.
- Wiener, N. (1948). *Cybernetics: Or control and communication in the animal and the machine*. MIT Press.
- Wolf, T., Debut, L., Sanh, V., Chaumond, J., Delangue, C., Moi, A., ... Rush, A. M. (2020). Transformers: State-of-the-art natural language processing. *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, 38–45.
- Zadeh, L. A. (1965). Fuzzy sets. *Information and Control*, 8(3), 338–353.

III. Power, Society, and Institutions

- Achille, M. (2019). *Necropolitics* (S. Corcoran, Trans.). Duke University Press.
- Appadurai, A. (1996). *Modernity at large: Cultural dimensions of globalization*. University of Minnesota Press.
- Arendt, H. (1958). *The human condition*. University of Chicago Press.
- Bentham, J. (1995). *The panopticon writings* (M. Božović, Ed.). Verso. (Original works published 1787–1791)
- Bijker, W. E., Hughes, T. P., & Pinch, T. J. (Eds.). (1987). *The social construction of technological systems: New directions in the sociology and history of technology*. Cambridge, MA: MIT Press.
- Castells, M. (1996). *The rise of the network society*. Blackwell.
- Collini, S. (2012). *What are universities for?* London: Penguin.
- Connell, R. (2007). *Southern theory: The global dynamics of knowledge in social science*. Cambridge: Polity Press.
- Dussel, E. (1985). *Philosophy of liberation*. Orbis Books.
- Fanon, F. (1963). *The wretched of the earth* (C. Farrington, Trans.). Grove Press. (Original work published 1961)

- Fanon, F. (2008). *Black skin, white masks* (R. Philcox, Trans.). Grove Press. (Original work published 1952)
- Foucault, M. (1995). *Discipline and punish: The birth of the prison* (A. Sheridan, Trans.). Vintage. (Original work published 1975)
- Giroux, H. A. (2014). *Neoliberalism's war on higher education*. Chicago: Haymarket Books.
- Habermas, J. (1984). *The theory of communicative action: Reason and the rationalization of society* (Vol. 1). Beacon Press.
- Hall, S. (1980). Encoding/decoding. In S. Hall, D. Hobson, A. Lowe, & P. Willis (Eds.), *Culture, media, language* (pp. 128–138). Hutchinson.
- Hobbes, T. (1996). *Leviathan* (R. Tuck, Ed.). Cambridge University Press. (Original work published 1651)
- Jasanoff, S. (2004). *States of knowledge: The co-production of science and social order*. London: Routledge.
- Jenkins, H. (2006). *Convergence culture: Where old and new media collide*. New York University Press.
- Latour, B. (2005). *Reassembling the social: An introduction to actor-network-theory*. Oxford: Oxford University Press.
- Locke, J. (1988). *Two treatises of government* (P. Laslett, Ed.). Cambridge University Press. (Original work published 1689)
- Lugones, M. (2008). The coloniality of gender. *Worlds & Knowledges Otherwise*, 2(2), 1–17.
- Machiavelli, N. (1998). *The prince* (H. C. Mansfield, Trans.). University of Chicago Press. (Original work published 1532)
- Marginson, S. (2016). *The dream is over: The crisis of Clark Kerr's California idea of higher education*. Oakland: University of California Press.
- Marx, K. (1990). *Capital: A critique of political economy* (Vol. 1, B. Fowkes, Trans.). Penguin. (Original work published 1867)
- McChesney, R. W. (2008). *The political economy of media: Enduring issues, emerging dilemmas*. Monthly Review Press.
- McLuhan, M. (1994). *Understanding media: The extensions of man*. MIT Press. (Original work published 1964)
- McMillan Cottom, T. (2017). *Lower ed: The troubling rise of for-profit colleges in the new economy*. New York: The New Press.
- Mill, J. S. (1978). *On liberty*. Hackett. (Original work published 1859)

Ngũgĩ wa Thiong'o. (1986). *Decolonising the mind: The politics of language in African literature*. Heinemann.

Pariser, E. (2011). *The filter bubble: What the Internet is hiding from you*. Penguin.

Pinch, T., & Bijker, W. E. (1984). The social construction of facts and artifacts: Or how the sociology of science and the sociology of technology might benefit each other. *Social Studies of Science*, 14(3), 399–441.

Quijano, A. (2000). Coloniality of power, Eurocentrism, and Latin America. *Nepantla: Views from South*, 1(3), 533–580.

Rawls, J. (1971). *A theory of justice*. Harvard University Press.

Readings, B. (1996). *The university in ruins*. Cambridge, MA: Harvard University Press.

Rodney, W. (1972). *How Europe underdeveloped Africa*. Bogle-L'Ouverture.

Rousseau, J.-J. (1997). *The social contract and other later political writings* (V. Gourevitch, Ed. & Trans.). Cambridge University Press. (Original work published 1762)

Said, E. W. (1978). *Orientalism*. Pantheon.

Santos, B. de S. (2014). *Epistemologies of the South: Justice against epistemicide*. Routledge.

Slaughter, S., & Rhoades, G. (2004). *Academic capitalism and the new economy: Markets, state, and higher education*. Baltimore: Johns Hopkins University Press.

Voltaire. (2009). *Candide and other stories* (R. Pearson, Trans.). Oxford University Press. (Original work published 1759)

Weber, M. (1978). *Economy and society: An outline of interpretive sociology* (G. Roth & C. Wittich, Eds.). University of California Press.

Winner, L. (1980). Do artifacts have politics? *Daedalus*, 109(1), 121–136.

Zuboff, S. (2019). *The age of surveillance capitalism: The fight for a human future at the new frontier of power*. PublicAffairs.

IV. Technology, Ecology, and Critical Futures

Amodei, D., Olah, C., Steinhardt, J., Christiano, P., Schulman, J., & Mané, D. (2016). Concrete problems in AI safety. *arXiv:1606.06565*.

Arthur, W. B. (1989). Competing technologies, increasing returns, and lock-in by historical events. *The Economic Journal*, 99(394), 116–131.

Arthur, W. B. (1994). *Increasing returns and path dependence in the economy*. University of Michigan Press.

- Arthur, W. B. (2009). *The nature of technology: What it is and how it evolves*. Free Press.
- Bacon, F. (2000). *The new organon* (L. Jardine & M. Silverthorne, Eds.). Cambridge University Press. (Original work published 1620)
- Brooks, R. (2002). *Robot: The future of flesh and machines*. Viking.
- Carey, J. W. (1989). *Communication as culture: Essays on media and society*. Routledge.
- Dator, J. (2009). Alternative futures at the Manoa School. *Journal of Futures Studies*, 14(2), 1–18.
- Ellul, J. (1964). *The technological society* (J. Wilkinson, Trans.). Vintage.
- Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. MIT Press.
- Haraway, D. J. (1991). *Simians, cyborgs, and women: The reinvention of nature*. Routledge.
- Heidegger, M. (1977). *The question concerning technology, and other essays* (W. Lovitt, Trans.). Harper & Row. (Original essay published 1954)
- Hinton, G. E., Osindero, S., & Teh, Y.-W. (2006). A fast learning algorithm for deep belief nets. *Neural Computation*, 18(7), 1527–1554.
- Inayatullah, S. (2004). The causal layered analysis (CLA) reader: Theory and case studies of an integrative and transformative methodology. Taipei: Tamkang University Press.
- Inayatullah, S. (2008). Six pillars: Futures thinking for transforming. *Foresight*, 10(1), 4–21.
- Jumper, J., Evans, R., Pritzel, A., Green, T., Figurnov, M., Tunyasuvunakool, K., ... Hassabis, D. (2021). Highly accurate protein structure prediction with AlphaFold. *Nature*, 596(7873), 583–589.
- McLuhan, M. (1994). *Understanding media: The extensions of man*. MIT Press. (Original work published 1964)
- Milojević, I. (2005). Educational futures: Dominant and contesting visions. London: Routledge.
- Morton, T. (2013). *Hyperobjects: Philosophy and ecology after the end of the world*. University of Minnesota Press.
- Ramos, J. M. (2019). Anticipatory governance: Traditions and trajectories for strategic design. *Journal of Futures Studies*, 23(4), 29–44.
- Russell, S. (2019). *Human compatible: Artificial intelligence and the problem of control*. Viking.
- Shiva, V. (1989). *Staying alive: Women, ecology and development*. Zed Books.
- Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., ... Hassabis, D. (2017). Mastering the game of Go without human knowledge. *Nature*, 550(7676), 354–359.

Wiener, N. (1948). *Cybernetics: Or control and communication in the animal and the machine*. MIT Press.

Wolf, T., Debut, L., Sanh, V., Chaumond, J., Delangue, C., Moi, A., ... Rush, A. M. (2020). Transformers: State-of-the-art natural language processing. *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, 38–45.

Technical Appendix: Evidence and Run Procedure

This appendix explains what sits in each evidence tier, where the data come from, and what the model does each time it runs. It is written for transparency without disclosing implementation details that are proprietary.

Tier 1 : Slow Anchors (frozen; refresh only on official releases)

Role. Structural context and denominators that change infrequently. We refresh Tier-1 only when the publisher posts a scheduled release that is past a short grace window.

Current Tier-1 indicators and sources

Macroeconomy

- **Real GDP growth (% y/y)** : Source: *IMF World Economic Outlook (WEO)*. Cadence: annual (Apr/Oct releases).
- **GDP per capita, PPP (constant int\$)** : Source: *World Bank World Development Indicators (WDI)*. Cadence: annual.

Population and human development

- **Total population (mid-year)** : Source: *UN DESA World Population Prospects (WPP)*. Cadence: annual/biennial.
- **Human Development Index (0–1)** : Source: *UNDP Human Development Report*. Cadence: annual.

Governance and institutional environment

- **Government Effectiveness (estimate)** : Source: *World Bank Worldwide Governance Indicators (WGI)*. Cadence: annual.
- **Corruption Perceptions Index (0–100)** : Source: *Transparency International (CPI)*. Cadence: annual.
- **Freedom in the World (aggregate scores)** : Source: *Freedom House*. Cadence: annual.
- **Mean institutional trust (0–100)** : Source: *Edelman Trust Barometer*. Cadence: annual (country coverage varies).

Energy structure

- **Electricity generation by fuel share (%)** : Source: *IEA Electricity Information*. Cadence: annual.

Digital access (structural)

- **Fixed-broadband subscriptions per 100 inhabitants** : Source: *ITU ICT Indicators*. Cadence: annual.

Education attainment and learning

- **Upper secondary completion rate (SDG 4.1.2)** : Source: *UNESCO Institute for Statistics (UIS)*. Cadence: annual.
- **PISA Creative Thinking (2022 mean score)** : Source: *OECD PISA 2022*. Cadence: quadrennial.

Global exposure (prices)

- **Food Price Index** : Source: *World Bank Commodities "Pink Sheet"*. Cadence: monthly (tracked as a structural exposure).

Tier 2 : Minimal Operational Signals (frozen set; vintages refreshed every run)

Role. High-frequency, public, and reproducible signals that capture service reliability and lived conditions. The list below is fixed. On each run we check each source for a new eligible vintage and incorporate it after basic validation.

Energy and grids

- **Nighttime-lights anomaly (index, 100=baseline)** : Source: *NOAA Earth Observation Group VIIRS (monthly)*. Global proxy for electricity service continuity.
- **Grid stress (operator module, monthly from sub-daily)** : Sources (where available): *ENTSO-E Transparency*; *US EIA-930*; *NGESO/BMRS (GB)*; *AEMO (Australia)*; *GRID-India/NLDC*; *Eskom system status (South Africa)*; *ONS SIN (Brazil)*. Metric: p95 forecast error or load-shedding stage-hours.

Health security

- **Influenza test positivity (%)** : Source: *WHO FluNet (weekly → monthly)*.
- **Excess mortality (z-score)** : Source: *EuroMOMO (weekly; participating European countries)*.

Education, skills, credentialing

- **Schools with Internet for pedagogical purposes, upper secondary (%), SDG 4.a.1** : Source: *UNESCO UIS (annual)*.
- **Teachers using AI "frequently/always" in instruction (%)** : Source: *OECD TALIS 2024 (cycle-based)*.
- **New DOIs per 12 months per million people** : Source: *Crossref (monthly rolling)*.
- **New ORCID registrations per 100k population** : Source: *ORCID (quarterly/latest available)*.

Trust and information integrity

- **Significant internet-outage events (count)** : Source: *CAIDA IODA (daily → monthly)*.
- **Confirmed national disruptions (count)** : Source: *NetBlocks (event-based → monthly)*, where coverage exists.

Treaty-lite coordination and cross-border flows

- **Container Throughput Index (index)** : Source: *RWI/ISL (monthly)*.
- **Liner Shipping Connectivity Index (index)** : Source: *UNCTAD (quarterly)*.

Compute and AI access

- **Fixed-broadband p95 latency (ms)** : Source: *Measurement Lab (M-Lab) NDT (monthly)*.
- **p10 download speed (Mbps)** : Source: *Measurement Lab (M-Lab) NDT (monthly)*.

Coverage notes: Some Tier-2 sources are global (e.g., NOAA, M-Lab). Others are region or country specific (e.g., ENTSO-E, EuroMOMO). Where an operator feed is unavailable, we rely on the global proxy in that family. Missing values are left missing rather than imputed.

Tier 3 : Interpretive and Experimentation Layer

Role. A focused application layer for countries, sectors, and geographies. Tier 3 turns the indicators above into lived narratives so readers can see what the numbers feel like. It is also where we trial recent public research or administrative sources that you provide.

Rules. Tier 3 never moves probabilities by itself. If Tier-3 work repeatedly reveals a measurable effect that is public, well defined, persistent, and non-redundant, we may promote it to Tier 2 in the next version.

What We Do On Each Run (plain language)

1. **Confirm the freeze.** Tier-0 canon and the Tier-1/Tier-2 lists are locked. No new sources are added mid-run.
2. **Search for updates.**
 - **Tier-1:** We check each publisher's release calendar. If there is a new official release past its grace window, we incorporate it. Otherwise Tier-1 remains unchanged.
 - **Tier-2:** For every indicator listed above, we query the source for a new eligible vintage. If present, we ingest it.
3. **Validate and align.** We apply basic checks (expected fields, sensible ranges, consistent units and time periods) and standardise units and time.
4. **Append vintages.** New values are stored alongside prior vintages so past runs are reproducible.
5. **Filter blips and corroborate.** We act only on signals that repeat across runs and that are corroborated by related indicators.
6. **Adjust within limits.** We make small, bounded probability nudges for the four outcomes only when the tests above pass. If a signal fades later, we walk the move back by the same step.
7. **Publish.** We release updated odds, a short "why it moved" note referencing the public indicator families and the new vintages, and the Tier-3 narratives. No private telemetry is used.

Disclosure scope (public vs proprietary)

We publish the architecture, the indicator **families** and their public sources, and the run procedure above. We reserve exact endpoints, transformation recipes, QA thresholds, and internal parameter values. Full materials can be provided to clients or independent auditors under NDA.

This appendix is intended to give an intelligent generalist a coherent picture of what the model is, what the data are, and what happens each time it runs, while preserving the proprietary details that constitute our IP.