

Access Without Displacement: An Access-Displacement Framework for AI Economic Transformation

Author: Vito Henjoto, Antecedent Labs

Affiliation: Antecedent Labs -- independent research company

ORCID: 0009-0003-0076-3665

© 2026 Vito Henjoto. All rights reserved. This work is licensed under the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License (CC BY-NC-ND 4.0). To view a copy of this license, visit <https://creativecommons.org/licenses/by-nc-nd/4.0/>

Abstract

For nearly a century, displacement predictions have consistently overestimated the labour market consequences of automation technologies. This paper proposes that the consistent directional error reflects a structural conflation: forecasters measure the speed of capability access and project displacement on a corresponding timescale, when the two dynamics are governed by different determinants and operate independently. The paper advances seven contributions. First, naming the unstated assumption dividing the displacement and complementarian traditions -- *AGI-R* (artificial general intelligence as replacement) versus *AGI-C* (artificial general intelligence as coexistence) -- transforms an implicit framing commitment into a testable empirical question, supported by five converging lines of evidence from biological neuroscience, philosophy of mind, deployment data, experimental studies, and meta-analysis. Second, the Access-Displacement Framework formalises the independence of capability access and economic displacement through two variables -- the access gap (G_A) and the displacement gap (G_D) -- with cross-generational evidence across five technology transitions spanning two centuries. Third, the Measurement Obsolescence Hypothesis documents accelerating degradation in Bureau of Labor Statistics benchmark data, with revisions escalating to 3.4 times the historical average in the absence of any declared recession. Fourth, the Sampling Residualisation Hypothesis identifies systematic survey nonresponse bias that disproportionately excludes workers who have adapted to artificial intelligence. Fifth, the Organisational Absorption Rate demonstrates that institutional integration capacity operates as a binding rate constraint regardless of technical capability -- a distinction absent from thirty-five years of absorptive capacity research. Sixth, the Economic Forcing Function shows that the deceleration of Moore's Law drives deployment toward hybrid architectures that structurally embed human participation. Seventh, the Hybrid Swarm Architecture describes the three-layer deployment topology these

economic constraints produce. Each contribution is presented as a testable hypothesis with specified falsification conditions, not as a conclusion.

Keywords: artificial intelligence, labour markets, displacement, augmentation, AGI, measurement, organisational absorption

1. Introduction

For nearly a century, each major advance in automation technology has produced a wave of predictions that human labour is about to become obsolete. Keynes (1930) coined the phrase "technological unemployment" while anticipating a fifteen-hour working week within a generation. Leontief (1983) warned that human workers would be displaced as decisively as horses were displaced by tractors. Frey and Osborne (2013) estimated that 47 per cent of United States occupations faced high risk of computerisation within ten to twenty years. Goldman Sachs (2023) identified 300 million jobs globally as "exposed" to generative artificial intelligence. Each prediction carried the authority of its moment -- a Nobel laureate, an Oxford working paper that would accumulate over 17,000 citations, a Wall Street research desk. Each responded to genuine technological change.

These predictions share a consistent structure. They estimate the technical potential for automation, project that potential as economic inevitability, and treat the resulting figures as displacement forecasts rather than exposure estimates. The distinction matters. Arntz, Gregory, and Zierahn (2016) demonstrated that reanalysing Frey and Osborne's data at the task level rather than the occupation level reduced the "at risk" figure from 47 per cent to 9 per cent across 21 OECD countries. The Organisation for Economic Co-operation and Development (2023), despite identifying 27 per cent of jobs across member states as at high risk of automation, concluded that there was "little evidence of significant job losses." The World Economic Forum (2025) projected a net gain of 78 million jobs by 2030. Forrester (2026) estimated that over half of layoffs attributed to artificial intelligence would be "quietly reversed" as organisations discovered operational challenges. Every major displacement prediction that has reached its forecast horizon has overestimated actual job loss. United States unemployment fell to 3.4 per cent in January 2023 -- the lowest rate since 1969 -- a decade after Frey and Osborne's initial publication.

The most extreme entry in this genre appeared in February 2026, when Citrini Research published "The 2028 Global Intelligence Crisis" (Citrini Research, 2026), projecting that artificial intelligence would drive the S&P 500 down 38 per cent, push unemployment to 10.2 per cent, and compress labour's share of gross domestic product from 56 per cent to 46 per cent -- all within two years. The essay introduced the concept of "Ghost GDP," in which productivity gains accrue to compute owners without circulating through the consumer economy, and modelled second-order financial contagion effects including private credit defaults and residential mortgage distress. The essay, which Citrini Research explicitly presented as a hypothetical scenario, used projected Bloomberg terminal

headlines as illustrative evidence. The Dow Jones Industrial Average fell approximately 800 points the following trading day; markets recovered within two days.

Citrini Research did not invent their analytical framework. Their occupation-level framing descends from Frey and Osborne (2013). Their "Ghost GDP" mechanism extends Frey's (2019) application of Allen's (2009) Engels' Pause -- the documented historical period during the first decades of British industrialisation in which productivity gains flowed to capital owners while real wages stagnated. Their investment-revenue gap builds on Cahn's (2024) "\$600 Billion Question" concerning the widening disparity between artificial intelligence infrastructure spending and actual revenue generation. Their timeline compression draws credibility from Amodei's (2024, 2026) characterisation of the human-AI "centaur phase" as "very brief." What distinguishes their essay from its predecessors is not the thesis but three features: a financial contagion mechanism absent from prior displacement literature, compression of the standard five-to-twenty-year forecast horizon to two years, and specific falsifiable financial targets.

Separately, Burry (2026) offered a more analytically grounded critique of artificial intelligence infrastructure investment, identifying two distinct claims: that major technology companies have systematically understated depreciation on artificial intelligence capital expenditure, and that the widening gap between infrastructure investment and revenue generation signals impending asset stranding. The first claim has found partial support in subsequent accounting revisions. The second rests on the same unstated assumption that pervades the displacement genre: that the current trajectory of centralised artificial intelligence deployment is the only trajectory, and that any deviation from it constitutes failure. This paper argues that the deviation is, in fact, the economically rational outcome.

The consistent directional error across nine decades of displacement prediction is not random. Five structural features recur across the genre. First, technical capability is treated as economic inevitability. Second, aggregate exposure figures are presented as displacement estimates -- what we term the *exposure-to-displacement slippage*. Third, historical precedent is selectively invoked: the Engels' Pause is cited while the Victorian employment boom that followed it is not. Fourth, adoption timelines are compressed or absent. Fifth, job creation is underspecified relative to job destruction. Acemoglu (2024), in the most rigorous quantitative assessment to date, demonstrated the cost of this final omission empirically: of the 19.9 per cent of total work tasks exposed to artificial intelligence, only 23 per cent were cost-effective to automate within a decade, yielding an effective automation rate of 4.6 per cent of all tasks, an order of magnitude below the headline "exposure" figures that institutional reports promote.

A substantial counter-tradition exists. Autor's (2003; 2015) task framework demonstrated that automation complements non-routine cognitive work and that the economy's capacity to generate new tasks has historically outpaced its capacity to automate existing ones. Bessen (2015) showed that automated teller machine deployment between 1985 and 2002 coincided with a net increase in bank teller employment, from approximately

485,000 to 527,000, even as the number of machines grew from 60,000 to 352,000. PwC (2018), in a rare institutional acknowledgment of adoption friction, explicitly discounted theoretical automation risk by two-thirds to account for "economic, legal, regulatory and organisational barriers to adoption" -- arriving at a net-neutral projection for the United Kingdom. Acemoglu, Autor, and Johnson (2026), in the most comprehensive recent treatment, proposed a five-category taxonomy of technological change -- distinguishing labour-augmenting, capital-augmenting, automating, expertise-levelling, and new-task-creating innovations. They identified three market failures that bias artificial intelligence deployment toward automation rather than augmentation. Autor and Thompson (2025), measuring expertise intensity across 303 harmonised occupations from 1980 to 2018, found that automation simultaneously raises wages in occupations where it eliminates inexpert tasks and lowers wages in occupations where it eliminates expert tasks. This dual effect resolves a long-standing puzzle in the polarisation literature and demonstrates that the consequences of automation are neither uniformly positive nor uniformly negative.

The experimental literature reinforces the complementarian position. Brynjolfsson, Li, and Raymond (2023) found that generative artificial intelligence produced a 13.8 per cent average productivity gain among customer support agents, with lower-skilled workers gaining approximately 35 per cent while top performers gained little or nothing. Dell'Acqua et al. (2023) documented a 40 per cent quality improvement for management consultants working within the technology's capability frontier, with below-average performers improving 43 per cent compared with 17 per cent for above-average performers; crucially, consultants working on tasks outside the frontier (beyond the technology's reliable capability) were 19 percentage points less likely to produce correct solutions compared with those working without artificial intelligence, demonstrating that the boundary between augmentation and degradation is domain-specific, not universal. This distributional compression, with lower performers gaining the most and higher performers gaining the least, is the single most robust empirical finding across all experimental studies of artificial intelligence augmentation (see also Noy and Zhang, 2023; Peng et al., 2023).

Acemoglu and Restrepo (2019), using Bureau of Labor Statistics data spanning 1947 to 2017, provided the empirical foundation that both traditions build upon. Between 1947 and 1987, displacement averaged approximately 0.48 per cent per year and new task creation approximately 0.47 per cent per year -- near-perfect balance for four decades. After 1987, this balance broke. Displacement rose modestly, from 0.48 to approximately 0.70 per cent per year. But reinstatement fell sharply, from 0.47 to approximately 0.35 per cent per year. The problem is not that displacement has accelerated dramatically; it is that new task creation has slowed. Acemoglu and Restrepo documented this structural shift but did not explain why reinstatement weakened. This paper offers candidate explanations.

Yet both traditions carry an unstated assumption. Every argument for human-AI complementarity implicitly presumes that artificial general intelligence, if and when it emerges, will coexist with human cognition rather than replace it. Every displacement prediction implicitly presumes the opposite. Neither tradition names this assumption. The

displacement literature assumes what we term *AGI-R* -- artificial general intelligence as replacement -- in which machine cognition substitutes for human cognition across all economically relevant dimensions. The complementarian literature assumes *AGI-C* -- artificial general intelligence as coexistence -- in which machine and human cognition remain structurally complementary regardless of machine capability. Despite more than seventy proposed definitions of artificial general intelligence (Legg and Hutter, 2007), none distinguishes between these two trajectories. Making this distinction explicit is this paper's first contribution, as it transforms an unstated framing assumption into a testable empirical question.

I argue that this consistent directional error is neither accidental nor evidence of analytical incompetence. It is the predictable result of a structural conflation. This paper makes seven contributions:

1. **AGI-C and AGI-R** name the unstated assumption underlying both the displacement and complementarian traditions, enabling their first direct empirical comparison. Five converging lines of evidence, from biological neuroscience, philosophy of mind, deployment data, experimental studies, and meta-analysis, support AGI-C as the dominant trajectory. Falsification conditions are specified.
2. The **Access-Displacement Framework** formalises the distinction between capability access and economic displacement, presenting cross-generational evidence that these dynamics are structurally independent. The framework introduces two formal variables -- the access gap (G_A) and the displacement gap (G_D) -- and demonstrates their independence across five technological generations spanning two centuries.
3. The **Measurement Obsolescence Hypothesis** documents accelerating degradation in the benchmark labour statistics on which displacement predictions depend. Bureau of Labor Statistics benchmark revisions have escalated from -187,000 in March 2023 (within historical norms) to -598,000 in March 2024 (2.3 times the historical average) to -862,000 in March 2025 (3.4 times the average) -- the largest final benchmark revision since the Great Recession.
4. The **Sampling Residualisation Hypothesis** identifies a systematic bias in survey nonresponse that disproportionately excludes workers who have adapted to artificial intelligence -- those most likely to work in non-traditional arrangements that establishment surveys are designed to miss.
5. The **Organisational Absorption Rate** demonstrates that institutional capacity to integrate new technology operates as a binding rate constraint regardless of technical capability -- a distinction absent from approximately 60,000 citations building on Cohen and Levinthal's (1990) absorptive capacity framework, which treats absorption as a stock rather than a flow.

6. The **Economic Forcing Function** shows that the deceleration of Moore's Law imposes escalating cost floors on centralised computation, making hybrid human-AI deployment the economically rational architecture rather than a transitional compromise.

7. The **Hybrid Swarm Architecture** describes the deployment pattern already emerging in production systems -- centralised training combined with distributed edge inference and bidirectional human-AI integration -- as the architectural expression of AGI-C under Economic Forcing Function constraints.

Together, these seven contributions constitute a unified framework for understanding why displacement predictions have been systematically wrong in the same direction for nearly a century, and what the empirical evidence reveals about the actual trajectory of artificial intelligence and labour interaction. Each contribution is presented as a testable hypothesis with specified falsification conditions, not as a conclusion.

The remainder of this paper proceeds as follows. Section 2 develops the Access-Displacement Framework and presents the historical evidence for the structural independence of access and displacement dynamics. Section 3 addresses the measurement crisis: the Measurement Obsolescence Hypothesis (Section 3.1) and the Sampling Residualisation Hypothesis (Section 3.2) demonstrate that the data substrate of displacement predictions is actively degrading. Section 4 identifies two absorption constraints: the Organisational Absorption Rate (Section 4.1) and the Economic Forcing Function (Section 4.2). Section 5 presents the coexistence architecture: AGI-C as the dominant trajectory (Section 5.1) and its architectural expression in Hybrid Swarm systems (Section 5.2). Section 6 discusses policy implications, cross-domain implications of skill compression, limitations, and directions for future research.

2. The Access-Displacement Framework

The displacement predictions surveyed in Section 1 share a structural feature that has received insufficient analytical attention. Each begins by estimating how quickly artificial intelligence capabilities have become accessible -- to firms, to workers, to entire economies -- and then projects economic displacement on a comparable timescale. The implicit reasoning is straightforward: if the technology is available, its labour market consequences must follow shortly. Citrini Research (2026) compress this to two years. Frey and Osborne (2013) allowed ten to twenty. Goldman Sachs (2023) left the timeline vague but implied urgency. In every case, the speed of access is treated as a reliable predictor of the speed of displacement.

This section introduces a framework that challenges that treatment. We define two formal variables -- the access gap and the displacement gap -- and present cross-generational evidence that they are structurally independent. If this independence claim holds, then any forecast that infers displacement timelines from access timelines commits a category error, regardless of how sophisticated its modelling of technical capability may be.

2.1 Formal Definitions

We define two temporal intervals that characterise the economic trajectory of any general-purpose technology.

The **access gap** (G_A) is the distance between available technological capability and current organisational utilisation. Operationally, it is the interval between a technology's commercial introduction and its widespread accessibility, measured as the time to majority adoption -- 50 per cent household or firm penetration, or an equivalent market threshold. The access gap is driven by production economics, distribution infrastructure, pricing dynamics, and network effects. It has compressed dramatically across successive technology generations: from approximately 40 years for the power loom and electrification to approximately two to three years for generative artificial intelligence. This compression is well documented. Comin and Hobijn (2010), analysing the cross-country adoption of over 100 technologies across 166 countries from 1820 to 2003, found that newer technologies are adopted faster and that cross-country adoption lags have converged. The access gap, in short, is an accelerating process.

The **displacement gap** (G_D) is the distance between current workforce levels and the level that would obtain if all technically feasible displacement were realised. Operationally, it is the interval between widespread accessibility and measurable economic restructuring -- the time from majority adoption to statistically significant employment decline in directly exposed occupations (greater than 10 per cent decline from peak, sustained for more than 24 months). The displacement gap is driven by a fundamentally different set of determinants: complementary innovation requirements (Bresnahan and Trajtenberg, 1995), organisational absorption capacity (Cohen and Levinthal, 1990), institutional adaptation, demand elasticity (Bessen, 2018), and labour market restructuring dynamics (Caselli, 1999). Where the displacement gap is measurable, it has remained in the range of 15 to 40 years regardless of how quickly the underlying technology became accessible. This invariance is the empirical puzzle that motivates the framework.

The **independence claim** is that G_A and G_D are structurally independent. The access gap can close without the displacement gap closing. Compression of the access gap does not predict compression of the displacement gap, because the two are driven by different determinants operating on different timescales. This is not merely the observation that displacement lags behind access, a point that David (1990), Brynjolfsson, Rock, and Syverson (2021), and Perez (2002) have each documented within individual technology cycles. It is the stronger claim that the *relationship itself* does not tighten across generations. As access has accelerated from decades to years, displacement has not followed.

This independence claim distinguishes the Access-Displacement Framework from three related bodies of work. David (1990) documented a 40-year lag between electricity adoption and factory productivity gains, but modelled it as a single coupled process with endogenous delay: complementary investments that the technology itself eventually

catalyses. Perez (2002) distinguished installation from deployment periods within 50-to-60-year techno-economic surges, but treated the two phases as causally linked through financial capital dynamics: installation-phase speculation provides the capital that eventually funds deployment. Fichman and Kemerer (1999) demonstrated at the firm level that technology acquisition and technology deployment follow distinct trajectories, with purchase-based diffusion metrics systematically overstating productive adoption. However, they made no independence claim and operated at the firm rather than the macroeconomic level. Each predecessor identifies a gap between access and impact. None claims that the gap is invariant to access speed across generations.

The framework yields a specific falsifiable prediction. If the access gap for generative artificial intelligence is two to three years, and if the independence claim holds, then measurable occupation-level displacement (greater than 10 per cent sustained decline in directly affected occupations) should not be observed before the late 2020s at the earliest. Forecasts projecting broad displacement by 2028, including Citrini Research (2026), would constitute a category error: projecting access-gap dynamics onto displacement-gap timescales.

2.2 Cross-Generational Evidence

The historical record provides five technology transitions for which both access and displacement dynamics can be assessed with sufficient data to evaluate the independence claim. We present them in chronological order.

The power loom (1785--1830). Edmund Cartwright patented the power loom in 1785, and factory-based weaving diffused across Lancashire and the English Midlands over the following four decades -- an access gap of approximately 40 years from patent to widespread factory adoption. The displacement consequences were severe for handloom weavers, whose numbers declined from approximately 250,000 in 1820 to near zero by 1860 (Allen, 2009). Yet the net effect on textile employment was the opposite of what a displacement-only analysis would predict. Factory employment in cotton textiles grew approximately fourfold between 1795 and 1835, absorbing displaced handloom weavers alongside new entrants at lower but rising real wages (Allen, 2009). Allen documented this period as the "Engels' Pause" -- approximately six decades (1770--1830) during which productivity gains from mechanisation flowed primarily to capital owners while real wages stagnated. The pause eventually broke, and real wages rose sharply after 1830. The power loom case thus exhibits the full pattern: rapid capability development, extended access gap, severe occupational displacement within the affected craft, but net employment growth in the broader sector -- with the displacement and access dynamics operating on visibly different timescales and responding to different determinants (factory economics, labour supply, demand for textiles, and the gradual development of complementary institutions including factory legislation and worker organisation).

Electrification (1880--1920). The electric motor became commercially available in the early 1880s and achieved widespread factory adoption over approximately 20 years. Yet

the economic restructuring it enabled -- the reorganisation of factory production from centrally powered shaft-driven layouts to unit-drive configurations -- required approximately 40 years. David (1990) documented that the productivity gains from electrification did not materialise until factories were physically redesigned around the new technology, a process that required complementary innovations in factory architecture, management practice, and workforce organisation that were themselves independent of electricity's availability. The access gap was approximately 20 years; the displacement gap -- measured as the time from widespread access to measurable productivity restructuring -- was approximately 40 years. Workers were redeployed within manufacturing rather than displaced from it, as the new factory configurations required different skills but not fewer workers.

Automated teller machines (1970--2010). The first ATM was installed at Chemical Bank in Rockville Centre, New York, on 2 September 1969 (Bessen, 2015). Nationwide coverage was achieved by the late 1990s -- an access gap of approximately 28 years. Between 1988 and 2004, the number of tellers per branch fell from approximately 20 to 13 as ATMs automated routine transactions. But the reduction in per-branch operating costs made new branches economically viable, and banks opened 43 per cent more urban branches over the same period (Bessen, 2015). Total teller employment rose from approximately 250,000 in 1970 to a peak of approximately 550,000 around 2007 -- a near-doubling over nearly four decades following ATM introduction. Displacement, when it arrived, was driven not by ATMs but by a successor technology wave: online banking, mobile deposits, and digital banking applications. Teller employment has since declined to approximately 347,400 by 2024, with the Bureau of Labor Statistics projecting a further 13 per cent reduction through 2034. The Burning Glass Institute (2025) documented that teller postings have fallen by nearly two-thirds since 2010 and that only 4 per cent of tellers transition to higher-paying roles -- undermining the "relationship banking" narrative as a long-term outcome.

The corrected ATM case strengthens rather than weakens the independence claim. Access was rapid; displacement did not occur for nearly four decades. When displacement did materialise, it was driven by a different technology entirely. The displacement gap was determined by organisational and technological dynamics independent of the original access gap. The causal drivers were not merely lagged -- they were different in kind.

Internet and e-commerce (1995--2015). Consumer Internet use reached majority household penetration by approximately 2005, roughly a decade after the first commercial web browsers. Amazon was founded in 1994. Measurable displacement of brick-and-mortar retail employment did not begin until approximately 2010, and significant restructuring -- characterised by sustained store closures, warehouse employment growth, and last-mile logistics expansion -- has unfolded over more than 15 years and remains ongoing. United States retail employment has declined less than 10 per cent from its peak despite e-commerce capturing over 15 per cent of total retail sales, because new roles in fulfilment, logistics, and digital commerce have partially offset traditional retail job losses.

The access gap was approximately 10 years; the displacement gap exceeds 20 years and is still in progress.

Generative artificial intelligence (2022--present). ChatGPT launched in November 2022. By early 2025, generative AI tools were accessible to the majority of knowledge workers in developed economies -- Crane, Green, and Soto (2025) estimated that 20 to 40 per cent of the United States workforce was actively using AI tools. The access gap has compressed to approximately two to three years, the fastest for any general-purpose technology in the historical record. The displacement gap has not yet begun to register in employment statistics. As Section 3 will demonstrate, the measurement infrastructure that would detect such displacement is itself degrading. Yet the absence of displacement signal even in the most advanced firm-level studies -- Brynjolfsson, Li, and Raymond (2023) found productivity gains but no net headcount reduction; Dell'Acqua et al. (2023) found augmentation within the capability frontier and degradation outside it, but no displacement in either case. These findings suggest that the displacement gap is, at minimum, far longer than the access gap.

Table 1. Access Gaps and Displacement Gaps Across Technology Generations

Technology	Period	Access Gap (G_A)	Displacement Gap (G_D)	G_D Outcome
Power loom	1785--1830	~40 years	Handloom weavers displaced	Textile employment grew ~4x (Allen, 2009)
Electrification	1880--1920	~20 years (David, 1990)	~40 years	Redeployment within manufacturing
ATMs / banking	1970--2010	~28 years	~38 years (different technology)	Teller numbers grew through 2007 (Bessen, 2015)
Internet / e-commerce	1995--2015	~10 years	~20+ years, ongoing	Retail displacement <10%; new roles in fulfilment
Generative AI	2022--present	~2--3 years	TBD	G_D minimal so far despite rapid G_A closure

Two patterns are visible across Table 1. First, where displacement is measurable, the displacement gap has consistently exceeded the access gap, typically by a factor of 1.5 to 2 or more. Second, the access gap has compressed dramatically across generations -- from 40 years for the power loom to two to three years for generative AI -- while displacement gaps show no corresponding compression, remaining in the range of 15 to 40 years regardless of access speed. If the historical pattern holds, generative AI's displacement gap should be measured in decades, not the two years that Citrini Research (2026) project.

We acknowledge the structural limitation of this evidence. Only five to six general-purpose technology transitions exist in the historical record. Statistical independence cannot be established from five observations. The framework is presented as a hypothesis with cross-generational illustrative evidence, not as a proven statistical relationship. Comin and Hobijn (2004; 2010), analysing the cross-country adoption of 104 technologies over two centuries, found that technology adoption speeds have converged across countries while economic restructuring patterns have not -- implying the independence of access and displacement dynamics without explicitly stating it. Rodrik (2016) documented "premature deindustrialisation" in developing economies: access to manufacturing technologies converged while the employment restructuring associated with industrialisation did not. These cross-country findings, while not identical to the Access-Displacement Framework's cross-generational comparison, provide supporting evidence from a substantially larger dataset and are consistent with the independence claim.

Five mechanisms, drawn from distinct literatures, explain why compression of the access gap does not entail compression of the displacement gap. Complementary innovation requirements mean that general-purpose technologies require downstream sectors to develop sector-specific adaptations before productive restructuring can occur (Bresnahan and Trajtenberg, 1995) -- and the timeline for such adaptations is determined by domain knowledge, regulatory environments, and institutional structures, none of which accelerate because the upstream technology has become more accessible. Demand elasticity determines whether automation increases or decreases employment in an affected sector (Bessen, 2018): when automation reduces unit costs and demand is price-elastic, output expands faster than labour savings, and total employment rises -- as in the ATM case. Skill composition constrains restructuring speed through the portability of the affected workforce's capabilities (Caselli, 1999). Organisational absorption operates as a rate constraint regardless of technical availability: firms' capacity to integrate new technology is limited by management bandwidth, process redesign capacity, and workforce retraining throughput (Cohen and Levinthal, 1990). And costly adoption creates an inverse short-run relationship: Greenwood and Yorukoglu (1997) demonstrated that faster technology adoption actually produces deeper short-run productivity drag, because simultaneous adoption across sectors creates overlapping learning costs and implementation failures -- the opposite of the positive correlation between access speed and displacement speed that crisis narratives assume.

2.3 The Turing Trap and the Reinstatement Problem

The Access-Displacement Framework intersects with a parallel distinction in the artificial intelligence literature. Brynjolfsson (2022), in his characterisation of the "Turing Trap," identified two trajectories for artificial general intelligence. *AGI-R* (replacement) describes a trajectory in which machine cognition substitutes for human cognition across all economically relevant dimensions -- the implicit assumption of the displacement tradition. *AGI-C* (coexistence) describes a trajectory in which machine and human cognition remain structurally complementary regardless of machine capability -- the implicit assumption of the complementarian tradition. Brynjolfsson argued that the research community's focus

on passing the Turing Test -- building machines that replicate human performance -- has inadvertently biased development toward AGI-R applications, when the greater economic value lies in AGI-C applications that augment rather than replace human capabilities. The trap, in Brynjolfsson's formulation, is not that replacement is inevitable; it is that replacement has become the default design objective.

The Access-Displacement Framework maps onto this distinction. Rapid G_A closure, the swift availability of powerful AI capabilities, is compatible with either trajectory. Under AGI-R, rapid access would predict rapid displacement: if the technology replaces human cognition, then access is the governing constraint, and displacement follows access. Under AGI-C, rapid access predicts rapid augmentation but not rapid displacement, because the technology complements human cognition rather than substituting for it, and the displacement gap remains governed by its own determinants: organisational absorption, complementary innovation, demand elasticity, and institutional adaptation. The historical evidence presented in Section 2.2 is consistent with AGI-C: access has accelerated across five technology generations without producing corresponding acceleration in displacement. Section 5 develops the AGI-C case in full, presenting converging evidence from biological neuroscience, philosophy of mind, deployment data, and experimental studies.

Acemoglu and Restrepo (2019) provided the empirical foundation that both trajectories build upon. Using Bureau of Labor Statistics data spanning 1947 to 2017, they demonstrated that between 1947 and 1987, displacement averaged approximately 0.48 per cent per year and new task creation (reinstatement) approximately 0.47 per cent per year -- near-perfect balance for four decades. After 1987, this balance broke. Displacement rose modestly, from 0.48 to approximately 0.70 per cent per year. But reinstatement fell sharply, from 0.47 to approximately 0.35 per cent per year. The problem that Acemoglu and Restrepo documented is not that displacement has accelerated dramatically; it is that new task creation has slowed. The deceleration is in reinstatement, not the acceleration of displacement. This reframes the policy problem entirely. If displacement were accelerating, the appropriate response would be displacement prevention -- regulation, taxation of automation, deceleration of adoption. But if the problem is that reinstatement has slowed, the appropriate response is new task creation: investment in the complementary innovations, institutional frameworks, and organisational capacities that generate new forms of productive human work. The Access-Displacement Framework identifies the mechanisms (complementary innovation, organisational absorption, demand elasticity) through which reinstatement operates. The Turing Trap identifies the design bias, the default orientation toward replacement rather than augmentation, that may explain why reinstatement has slowed. Together, they suggest that the operative constraint on labour market adjustment to artificial intelligence is not technical capability but institutional and organisational capacity to create new tasks that exploit human-AI complementarity. Naming the AGI-C and AGI-R trajectories explicitly (a contribution developed in Section 5) transforms this from an unstated framing assumption into a testable empirical question.

2.4 Implications for Measurement and Prediction

The Access-Displacement Framework, if the independence claim holds, carries three implications for how artificial intelligence's labour market consequences should be analysed.

First, displacement predictions built on access-gap indicators are systematically unreliable. Adoption rates, user counts, capability benchmarks, and infrastructure investment figures measure G_A , not G_D . Projecting displacement timelines from these indicators -- as Frey and Osborne (2013), Goldman Sachs (2023), and Citrini Research (2026) each do, with varying degrees of explicitness -- imports the access gap's dynamics into a domain governed by different determinants. The consistent directional error of displacement predictions documented in Section 1 is precisely what the framework predicts: forecasters measure access, observe acceleration, and infer displacement on a corresponding timescale. The displacement gap, governed by organisational absorption and complementary innovation rather than by technical availability, does not cooperate.

Second, measurement systems built on the assumption that access and displacement are coupled will produce systematically misleading signals. If the standard labour market surveys were designed to detect displacement patterns characteristic of previous technology transitions -- gradual, sector-specific, observable through establishment-level headcount changes -- and if the current transition is instead characterised by rapid access with delayed and structurally different displacement dynamics, then the measurement infrastructure itself becomes a source of error. Section 3 develops this argument in detail, demonstrating that the Bureau of Labor Statistics benchmark revisions have escalated from historical norms to unprecedented magnitudes precisely during the period of fastest AI access-gap closure -- and that this escalation is consistent with a measurement apparatus that cannot distinguish between "no displacement" and "displacement that has taken forms the surveys were not designed to capture."

Third, the framework redirects analytical attention from the question that dominates the displacement literature -- "how many jobs will AI eliminate?" -- to the question that the evidence suggests is more tractable and more policy-relevant: "what determines the rate at which organisations create new tasks that exploit human-AI complementarity?" The access gap is closing rapidly. That much is not in dispute. The displacement gap has not yet opened. The analytical challenge is not to predict when it will, but to understand what governs the rate of reinstatement -- the creation of new productive tasks -- and whether that rate can be accelerated through institutional design, investment, and policy. Acemoglu and Restrepo (2019) documented that reinstatement slowed after 1987 without explaining why. The mechanisms identified in Section 2.2 -- complementary innovation, organisational absorption, demand elasticity, skill composition, and costly adoption dynamics -- offer candidate explanations. Sections 4 and 5 develop two of these mechanisms in detail: the Organisational Absorption Rate and the Economic Forcing Function.

3. The Measurement Crisis

Section 2 demonstrated that the access gap for generative artificial intelligence has compressed to approximately two to three years while the displacement gap -- governed by organisational absorption, complementary innovation, and institutional adaptation -- has not yet opened. Every displacement prediction cited in Section 1 depends, however, on a further assumption: that the labour market data against which predictions are calibrated accurately represents the economy in which displacement is supposed to occur. This section examines that assumption and finds it increasingly untenable.

The Bureau of Labor Statistics operates two primary employment measurement systems. The Current Employment Statistics (CES) survey samples approximately 670,000 establishment worksites against a universe of approximately 10 million employer establishments drawn from the Quarterly Census of Employment and Wages (QCEW). The Current Population Survey (CPS) contacts approximately 60,000 households monthly via telephone and in-person interviews. Both systems were architected for a labour market in which the dominant employment relationship is a W-2 payroll position with a single identifiable employer. Two independent lines of evidence -- escalating benchmark revisions and systematic survey nonresponse bias -- demonstrate that this measurement infrastructure is actively degrading. Crisis forecasts that build on this data inherit its distortions.

3.1 The Measurement Obsolescence Hypothesis

The CES benchmark process provides a natural diagnostic for measurement system health. Each year, the Bureau of Labor Statistics compares its monthly survey-based employment estimates against comprehensive counts from the QCEW -- an administrative dataset constructed from state unemployment insurance tax records that covers approximately 97 per cent of nonfarm wage and salary employment. The difference between survey estimates and QCEW counts, expressed as the benchmark revision, reflects the accumulated estimation error over the preceding twelve months. Benchmark revisions are expected; the historical average absolute revision for the period 2002 to 2023 is approximately 255,000 (Congressional Research Service, 2025). What is not expected is the pattern that has emerged since 2024.

Table 2. CES Benchmark Revision Escalation, 2023--2025

Benchmark Year	Preliminary Revision	Final Revision	Multiple of 2002--2023 Mean
March 2023	-306,000	-187,000	0.7x (within historical norms)
March 2024	-818,000	-598,000	2.3x
March 2025	-911,000	-862,000	3.4x

Sources: BLS CES National Benchmark Article; Congressional Research Service Report IF12827; BLS Preliminary Benchmark press releases.

Three features of this pattern warrant attention. First, the March 2023 revision was within historical norms -- the escalation begins with the March 2024 cycle, producing two consecutive outsized revisions rather than three. Second, the March 2025 preliminary revision of -911,000 briefly exceeded the previous all-time record of -902,000 set during the March 2009 benchmark cycle at the depth of the Great Recession; the final figure of -862,000, while slightly below that record, remains 3.4 times the historical average in the absence of any declared recession. Third, all three recent revisions are downward: initial estimates systematically overstated employment. Under a null hypothesis of unbiased estimation, three consecutive downward revisions has a prior probability of 12.5 per cent. Two consecutive revisions exceeding twice the historical mean, both downward, is considerably less likely under a stationary noise process.

The **Measurement Obsolescence Hypothesis (MOH)** proposes that this escalating revision pattern is not transient estimation error but a structural signal: the CES measurement framework is becoming progressively incompatible with the economy it is designed to measure. The mechanism operates through the birth-death model, which supplements the CES sample by estimating net job creation from business formations and closures that the survey cannot capture in real time. The birth-death model's estimates are calibrated against historical patterns of employer establishment formation and closure. When the structure of economic activity changes -- specifically, when traditional W-2 payroll employment gives way to 1099 independent contracting, platform-mediated work, and AI-augmented micro-enterprises -- the historical calibration becomes systematically wrong. The model overestimates business births (because new economic activity does not register as employer establishments in unemployment insurance records) and underestimates the displacement of traditional firms by structures the sampling frame does not capture. The birth-death model has systematically overestimated net business births for three consecutive years (2023 to 2025), and in January 2026 the Bureau of Labor Statistics modified the model to incorporate current sample information monthly rather than relying solely on historical patterns -- an institutional adaptation that the hypothesis predicts but that may prove insufficient if the structural gap continues to widen.

The CES-CPS divergence provides corroborating evidence. The establishment survey showed year-over-year employment growth of approximately 1.9 per cent in early 2025; the household survey showed approximately 0.4 per cent -- the smallest growth outside the pandemic period since October 2013. This divergence is precisely what structural measurement failure would produce. The CES counts payroll positions and cannot capture self-employed individuals or independent contractors. The CPS attempts to count all employed persons but relies on household surveys with response rates that have declined from approximately 82 per cent in 2019 to a series low of 64 per cent in November 2025. If structural transformation is moving workers from payroll to non-traditional arrangements, the CES would overcount (because the birth-death model overestimates traditional job creation) while the CPS would undercount (because non-traditional workers are harder to

reach and harder to classify). Two measurement systems designed for the same economy, both failing in opposite directions, is the signature of measurement framework obsolescence rather than random estimation error.

The institutional disruption timeline compounds the structural problem without being its cause. The Bureau of Labor Statistics has experienced a 22 per cent real decline in funding since fiscal year 2010 (Center on Budget and Policy Priorities). In August 2025, the BLS Commissioner was dismissed following a weak employment report. A 43-day government shutdown from October to November 2025 suspended all CPS data collection, rendering October household survey data unrecoverable. A second shutdown in January 2026 delayed the January employment report. The proposed fiscal year 2026 budget includes a further 8 per cent nominal cut to BLS funding and a 17 per cent reduction in full-time equivalent staffing. These institutional disruptions degrade measurement *capacity*; the Measurement Obsolescence Hypothesis addresses a distinct problem: measurement *framework* failure. Even a fully funded Bureau of Labor Statistics operating the current CES and CPS architecture would increasingly miss an economy structured around AI-augmented independent work, because the measurement framework was designed for a different economic structure. The institutional degradation accelerates the failure and reduces the Bureau's capacity to adapt, but the structural obsolescence precedes it. The March 2025 benchmark revision was calculated from QCEW data that predates the commissioner dismissal and government shutdowns.

The COVID-19 pandemic provides a critical natural experiment for disambiguating the current pattern from post-pandemic measurement recovery. A reviewer might reasonably argue that the 2024 and 2025 revisions reflect residual pandemic distortion -- that the CES birth-death model, its seasonal adjustment factors, and its population controls are still recovering from the unprecedented disruption of 2020-2021. Three features of the evidence argue against this interpretation. First, the *direction* of the current revisions is inconsistent with pandemic recovery noise. If the current revisions were pandemic aftershock, they should be distributed around zero -- sometimes overstating, sometimes understating -- as the measurement system corrects oscillating errors. Instead, all recent revisions are unidirectional: the CES has systematically overstated employment in every benchmark cycle since 2023. Second, the *acceleration* is inconsistent with recovery dynamics. Pandemic-induced measurement error should attenuate over time as statistical models recalibrate and response patterns normalise. Instead, the revisions are escalating: -187,000 in March 2023 (within historical norms), -598,000 in March 2024, -862,000 in March 2025. If the driving force were residual pandemic distortion, the March 2023 revision -- closest to the pandemic -- should have been the largest, not the smallest. Third, the COVID-era benchmark itself provides a direct comparison. The March 2021 annual benchmark revision was -7,000: the smallest in the 2002 to 2023 series, and effectively zero. Monthly CES estimation errors during 2020-2021 were large -- in-person data collection was suspended, seasonal adjustment factors were distorted, and employment patterns were historically unprecedented. But the annual benchmark correction was negligible because the underlying economic structures remained intact. When the economy reopened, payroll employment returned to forms the CES was designed to

capture. The measurement framework was stressed but not obsolete. The current period exhibits the opposite pattern: moderate monthly estimation variance, but benchmark-level corrections escalating year over year. No exogenous shock explains the escalation. The pattern is that of an estimation framework whose structural assumptions are diverging from economic reality at an accelerating rate, not a system recovering from a transient disruption.

This analysis yields a specific falsification condition. The March 2026 benchmark revision, expected in preliminary form by approximately August 2026, will provide the confirmatory test. If the March 2026 revision returns to within one standard deviation of the 2002-2023 historical mean, the MOH is substantially weakened: the 2024-2025 escalation would more plausibly reflect transient factors, including residual pandemic adjustment, rather than structural measurement obsolescence. If the March 2026 revision continues the escalating pattern -- a third consecutive outsized downward revision -- the MOH gains significant predictive credibility, because three consecutive anomalous revisions of increasing magnitude, all in the same direction, cannot be parsimoniously explained by transient noise or pandemic recovery.

Leduc, Oliveira, and Paulson (2025) concluded in an SF Fed Economic Letter published in March 2025 that declining survey response rates had not yet produced outsized revision volatility. We do not dispute this finding for the period they examined. Their analysis, however, was completed before the -911,000 preliminary benchmark revision (September 2025), before the commissioner dismissal (August 2025), before the 43-day government shutdown (October-November 2025), and before the CPS response rate reached its series low of 64 per cent (November 2025). More fundamentally, their paper tests for increased *noise* -- whether revision volatility has risen. The Measurement Obsolescence Hypothesis predicts directional *bias*: systematic overstatement of employment with increasing magnitude. A stationary noise process does not produce escalating unidirectional revisions. The distinction between noise (increased variance around zero) and bias (systematic directional drift) separates a measurement system that is merely noisier from one that is structurally misaligned.

Two recent contributions demonstrate that the measurement challenge identified by the MOH is not merely theoretical. Soto (2025), in a Federal Reserve Board Finance and Economics Discussion Series paper, constructed the AI Research Index (AIR Index) by extracting artificial intelligence research and development investment signals from corporate earnings call transcripts. The AIR Index captures dynamics in firm-level AI adoption that establishment surveys cannot detect -- it measures commitment intensity, investment trajectory, and strategic orientation toward AI integration through natural language signals that exist entirely outside the CES sampling frame. The AIR Index does not replace traditional measurement, but it demonstrates operationally that non-traditional measurement instruments can capture economically significant dynamics that establishment surveys miss. Crane, Green, and Soto (2025), in a Federal Reserve Board FDS Notes publication, estimated that 20 to 40 per cent of the United States workforce was actively using AI tools in the workplace -- a figure derived from survey instruments

distinct from the CES and CPS. Together, these contributions constitute proof of concept for what we term **Novel Temporal Labour Indicators (NTLI)**: measurement approaches that track economic transformation through behavioural, transactional, and linguistic signals rather than through establishment headcounts and household survey responses. The MOH does not merely identify a problem; it motivates the development of measurement architectures suited to the economy that is emerging rather than the economy that existed when current surveys were designed.

3.2 The Sampling Residualisation Hypothesis

The Measurement Obsolescence Hypothesis addresses the establishment survey's structural failure. The Sampling Residualisation Hypothesis (SRH) addresses a complementary failure in the household survey: the systematic exclusion from measurement of precisely the population whose economic behaviour is most informative about AI-driven transformation.

The Current Population Survey's response rate has declined from approximately 82 per cent in 2019 to a series low of 64 per cent in November 2025 for the basic monthly survey. The Annual Social and Economic Supplement response rate has declined from 69 per cent to 62 per cent over the same period (Census Bureau, 2025). Response rate decline alone is not informative; all household surveys have experienced declining participation. What is informative is *who* has stopped responding and *how* that systematic nonresponse correlates with the economic transformation the survey is designed to measure.

Three findings establish the mechanism. First, Bernhardt, Munro, and Wolcott (2024), published in the *Journal of Applied Econometrics*, documented that the share of households refusing to participate in the CPS tripled within a decade, and that nonresponse bias accounts for at least 10 per cent of the reported decline in the labour force participation rate between 2000 and 2020. Rising refusals systematically suppress measured employment and labour force participation. Second, the Census Bureau's annual evaluations of CPS nonresponse, which link nonrespondents to Internal Revenue Service W-2 records and Social Security Administration data, found that since 2020, CPS respondents have systematically *higher* W-2 earnings than the full population -- a reversal of the pre-2020 pattern in which the gap was negligible. By 2025, upward bias in Hispanic household income reached 3.8 per cent (Census Bureau, 2025). This reversal indicates a compositional change in who responds: the remaining respondent pool is disproportionately composed of traditional W-2 workers, while higher-earning individuals in non-traditional arrangements are systematically absent. Third, Abraham et al. (2021), published in the *Journal of Labor Economics*, demonstrated that a large majority of workers classified as self-employed in IRS administrative records are not classified as self-employed in the CPS -- a measurement gap that reflects both proxy respondent effects and respondent self-classification. Bracha and Burke (2023), in a Federal Reserve Bank of Boston working paper, found that CPS employment rates were "consistently understated from 2015 through 2022 due to the misclassification of gig workers as either unemployed or not in the labor force," estimating that the employment-to-population ratio would be

0.25 to 1.1 percentage points higher under conservative assumptions, and as much as 5.1 percentage points higher under generous assumptions. Their companion piece noted that "as many as 7 million gig workers aren't being counted" and, critically, that "gig workers sometimes aren't counting themselves."

The Sampling Residualisation Hypothesis proposes that these findings are connected by a self-reinforcing mechanism. AI tools become accessible. Workers who adopt them become more productive in independent settings. Productive workers migrate from W-2 payroll to independent contracting, freelance platforms, and micro-enterprise structures. The CPS survey design -- binary employed/unemployed classification via household telephone and in-person interviews -- fails to capture this population accurately. These workers are less likely to respond to the CPS because their economic status is precisely the variable correlated with nonresponse: they are Missing Not At Random (Rubin, 1976). The measurement gap widens as the phenomenon it should capture accelerates. This is not static nonresponse bias, in which a fixed proportion of the population is hard to reach. It is dynamic: the hard-to-reach population grows in proportion to the economic transformation being measured.

The scale of the unmeasured population is substantial. Estimates of independent work range from 58 million -- 36 per cent of employed respondents -- (McKinsey, 2022, survey of 25,062 adults) to 72.9 million (MBO Partners, 2025), depending on definitional scope. Goldman Sachs (Rindels, 2025) calculated that if gig and platform work were fully captured, the employment-to-population ratio would be approximately 65 per cent rather than the measured 60 per cent -- a five-percentage-point gap. MBO Partners (2025) found that 74 per cent of independent workers now use generative artificial intelligence, saving an average of nine hours per week. The AI-independent work connection is not hypothetical: it is documented at the population level, and it identifies a workforce segment that is simultaneously the most AI-adapted and the least visible to the CPS.

A diagnostic signature distinguishes workforce migration from workforce destruction. If declining payroll employment reflected demand destruction, wages would fall with employment -- employers cutting positions because demand has contracted. If declining payroll employment reflects migration to higher-value independent work, wages among payroll stayers should remain stable or rise. The ADP Research Institute has documented a "continuous and dramatic slowdown" in payroll job creation over three years alongside stable wage growth of 4.4 to 4.6 per cent for job stayers. This payroll-wage decoupling signature is consistent with the migration interpretation and inconsistent with the destruction interpretation.

The empirical test architecture for the SRH requires data that is available but not routinely published. The core test would regress state-level CPS nonresponse rates (available within Census Bureau systems but requiring FOIA or direct collaboration) against state-level IRS Schedule C filing growth, Census Nonemployer Statistics, and QCEW payroll employment trends, controlling for demographics, industry composition, broadband penetration, COVID severity, and immigration flows. The predicted result is that states with the largest

growth in independent work show the largest increases in CPS nonresponse. Ward and Edwards (2021), published in *Labour Economics*, found that geographic variation in COVID-era nonresponse was not correlated with COVID case rates, eliminating the most obvious alternative explanation and opening space for the structural mechanism the SRH proposes. A complementary test, extending Abraham et al.'s (2021) methodology, would examine whether the gap between IRS self-employment and CPS self-employment is growing faster in AI-exposed sectors -- professional services, information technology, creative industries -- than in sectors with lower AI adoption. If confirmed, this would directly support the SRH: AI-augmented workers migrate to self-employment structures that the CPS fails to capture.

3.3 Combined Implications

The Measurement Obsolescence Hypothesis and the Sampling Residualisation Hypothesis describe the same structural phenomenon from different measurement perspectives. The MOH addresses the establishment survey: the CES overestimates traditional employment because its birth-death model assumes historical business formation patterns that no longer hold. The SRH addresses the household survey: the CPS underestimates non-traditional employment because workers who have adapted to AI are systematically absent from the respondent pool. Together, they produce a double distortion: conventional labour force data overstates the population available to be displaced while understating the population that has already adapted. Any displacement forecast calibrated against this data inherits both biases.

This double distortion interacts directly with the Access-Displacement Framework developed in Section 2. If measurement systems were designed for an economy in which access and displacement are tightly coupled -- in which technological adoption produces proportional employment restructuring -- then those systems will interpret rapid access-gap closure as evidence of imminent displacement, because that is the only template they possess. But the independence claim established in Section 2 holds that access and displacement are structurally independent. A measurement infrastructure that cannot distinguish between "workers displaced from employment" and "workers who have migrated to employment arrangements the survey does not capture" will systematically overstate displacement and understate adaptation. The measurement crisis does not merely reduce statistical precision. It introduces a directional bias that reinforces the displacement narrative: every worker who leaves a W-2 payroll position for AI-augmented independent work registers as a job lost in the CES and, if they do not respond to the CPS, as a worker who has left the labour force entirely. The measurement apparatus converts adaptation into displacement by design.

The policy implications of navigating with degrading instruments are not academic. Federal Reserve interest rate decisions, fiscal stimulus calibration, unemployment insurance design, and workforce development investment all depend on labour force statistics that the MOH and SRH suggest are systematically distorted. If the employment-to-population ratio is five percentage points higher than measured (Goldman Sachs, Rindels, 2025), the

economy is substantially tighter than official data indicates. Monetary policy calibrated to understated employment data risks being too accommodative; fiscal policy risks misallocating resources toward job creation programmes in sectors where the bottleneck is not demand for workers but the measurement system's inability to see workers who already exist.

The combined argument does not claim that displacement is impossible or that labour market disruption is illusory. It claims that the empirical foundation on which displacement predictions rest is actively degrading, and that the direction of degradation (overstating traditional employment, understating adaptive employment) systematically inflates displacement estimates. Even if we accept the full force of this argument, however, a deeper question remains. If measurement is fixed -- if we could see all workers, in all arrangements, in real time, would displacement still fail to materialise at the rates the crisis literature predicts? Section 4 argues that it would, because two independent constraints, organisational absorption capacity and the economics of computation, limit the rate at which AI can restructure labour markets regardless of what measurement systems report.

4. Absorption Constraints

Even if the measurement crisis documented in Section 3 were resolved -- if benchmark revisions returned to historical norms and household surveys captured the full spectrum of employment arrangements -- two independent constraints would still prevent displacement from proceeding at the speed crisis narratives assume. The first is organisational: institutions cannot absorb transformative technology faster than their internal complexity allows, regardless of the technology's capability. The second is economic: the cost structure of artificial intelligence deployment itself drives architectural evolution toward coexistence rather than replacement. This section examines each.

4.1 Organisational Absorption Rate

Displacement predictions implicitly assume that adoption follows automatically from capability. Once artificial intelligence can perform a task, the logic runs, the organisation will deploy it and the worker who performed that task will be displaced. JPMorgan Chase can access the same large language models as a three-person startup. The question is not whether organisations *can* adopt artificial intelligence but how *fast* they can integrate it into their operations, and what happens to their workforce during the integration. The binding constraint is rate, not capacity. It is a flow variable, not a stock variable. This distinction has been absent from the literature for thirty-five years.

The absorptive capacity literature, originating with Cohen and Levinthal's (1990) foundational paper (approximately 60,000 citations, one of the most-cited papers in management science), addresses whether and how organisations absorb external knowledge. Zahra and George (2002, approximately 15,000 citations) distinguished potential from realised absorptive capacity, introducing a typology that implicitly contains a temporal element (potential must precede realisation) without formalising the speed of

conversion as a variable. Lane, Koka, and Pathak (2006, approximately 5,000 citations), reviewing 289 papers across 14 journals, came closest to the gap this paper identifies: they explicitly noted "time delays in the absorptive capacity process, accrued from gathering external data, using it in R&D, and utilizing it in production." But they did not measure those delays, did not connect them to workforce outcomes, and did not use them as constraints on macroeconomic displacement timelines. Across 35 years of research, the absorptive capacity tradition treats absorption as a stock variable (a property of the organisation) rather than a flow variable (the rate at which that property converts to structural change). The distinction is analogous to the difference between a container's volume and the rate at which it fills.

We introduce the **Organisational Absorption Rate (OAR)** as the flow-variable complement to absorptive capacity: the speed at which an organisation converts artificial intelligence capability into workforce structural change, and the macroeconomic implications of that rate constraint. No prior work in the absorptive capacity tradition formalises the absorption rate as a measurable variable or uses it to constrain macroeconomic displacement timelines.

The OAR argument rests on an upper bound, not a statistical sample. JPMorgan Chase represents the ceiling of institutional artificial intelligence absorption. In 2025, the bank allocated approximately \$2 billion specifically to artificial intelligence within a total technology budget of approximately \$20 billion and reported approximately \$2 billion in AI-derived benefits -- effectively 100 per cent return on AI-specific investment (Dimon, Bloomberg, October 2025; JPMorgan 2026 Company Update, 23 February 2026). Separately, eFinancialCareers (Butcher, 23 February 2026) reported \$150 million in coding efficiencies within the technology team specifically -- a narrow component of the bank-wide return, not a bank-wide figure. The institution satisfies every condition that should maximise absorption speed: maximum CEO commitment (Dimon publicly describes the bank as "fundamentally rewired"), unlimited capital, explicit willingness to restructure ("fewer jobs in certain functions," with "huge redeployment plans"), and access to the best available technical talent.

Despite these conditions, the bank's workforce outcome is reallocation, not displacement. Operations headcount declined approximately 4 per cent. Support functions declined approximately 2 per cent. Revenue-facing headcount grew approximately 4 per cent. Technology headcount grew approximately 1 per cent. Net headcount change was approximately zero. Operations became 6 per cent more efficient per employee. Fraud costs fell 11 per cent per unit. Software engineer productivity rose 10 per cent. The productivity gains are real; the workforce effect is redistribution, not reduction.

Two variables that displacement predictions conflate are, at JPMorgan, empirically independent: financial return on artificial intelligence investment and workforce displacement. The bank achieved approximately 100 per cent return on \$2 billion in AI-specific investment through operational efficiency gains and internal labour reallocation, not through net headcount reduction. This decoupling is predicted by the OAR framework:

organisations capture artificial intelligence value through redistribution because the organisational absorption rate constrains the speed of structural change. High return on investment does not imply, and does not produce, proportional displacement. They are independent variables.

If the institution with the largest technology budget in global banking -- one that has achieved demonstrable return and whose CEO publicly champions aggressive adoption -- produces redeployment rather than displacement, then the economy as a whole cannot displace faster than JPMorgan's rate implies. JPMorgan is the ceiling. The average institution is below it. A reviewer who objects that "one firm is not a sufficient sample" misunderstands the logic. This is an upper bound argument, not a sampling argument. You do not need a sample of Olympic sprinters to know the average human cannot run 100 metres in 9.58 seconds. You need one data point: the fastest.

The Klarna reversal provides the strongest independent confirmation -- and eliminates the most common counter-argument, that JPMorgan's flat headcount reflects regulatory constraints specific to banking. Klarna is a fintech company: small (approximately 3,400 employees after cuts, versus JPMorgan's 318,000), agile (modern technology stack, no legacy systems), minimally regulated (no bank charter, no prudential supervision at JPMorgan's level), and led by a CEO maximally committed to artificial intelligence displacement. Between 2022 and 2024, Klarna cut from approximately 7,000 to approximately 3,400 employees and deployed an artificial intelligence chatbot handling 75 per cent of customer interactions across 35 languages. In February 2024, CEO Sebastian Siemiatkowski claimed "AI is already doing the jobs humans used to do." By mid-2025, Klarna reversed course. Customer satisfaction had declined. Operational quality had degraded. Siemiatkowski admitted the company had "focused too much on efficiency and cost" and that "the result was lower quality, and that's not sustainable" (Bloomberg, 8 May 2025). The company began rehiring human agents.

What makes Klarna analytically decisive is not that the artificial intelligence failed. It did not. The chatbot handled 75 per cent of interactions. The technology performed within its capability frontier. What failed was organisational integration. The institution could not absorb the consequences of its own displacement strategy: the downstream effects on service quality, customer relationships, and operational coherence exceeded the organisation's capacity to manage them at the speed the strategy demanded. This is precisely the distinction OAR formalises. Absorptive capacity asks whether the organisation can adopt the technology (Klarna could, and did). OAR asks how fast the organisation can restructure around it without systemic degradation. Klarna's answer was: not as fast as its CEO intended. The constraint is organisational, not regulatory, not technological, and not a function of insufficient commitment. Removing every advantage the displacement literature assumes should accelerate adoption -- removing regulation, removing legacy systems, removing scale -- did not remove the rate constraint.

Cross-validation from other institutions confirms the pattern without being required for the upper bound argument. Wells Fargo CEO Charles Scharf stated in 2026 that the bank had

"not reduced headcount because of AI so far." Goldman Sachs projects a 4 per cent headcount reduction -- a projection, not an achievement, and at a rate inconsistent with macroeconomic crisis within two years.

Enterprise-level survey data generalises the finding beyond individual institutions. McKinsey (2025) reports that 88 per cent of organisations have adopted artificial intelligence but only 7 per cent have fully scaled it -- an 81 percentage-point gap that is the aggregate expression of OAR operating at population scale. S&P Global (2025) reports that 42 per cent of companies have abandoned most artificial intelligence initiatives, up from 17 per cent in 2024 -- abandonment is increasing, not decreasing, as early-stage enthusiasm confronts implementation reality. MIT's Project NANDA (Challapally et al., 2025) finds that 95 per cent of custom enterprise generative artificial intelligence projects fail to reach production or deliver measurable return on investment. Deloitte (2026), in "The State of AI in the Enterprise: The Untapped Edge" -- surveying 3,235 respondents across 24 countries and 6 industries -- reports that only 21 per cent of companies planning agentic artificial intelligence have mature governance frameworks in place. These findings are not evidence of technological failure. They are evidence of an organisational rate constraint operating at population scale: firms can access artificial intelligence (high absorptive capacity) but cannot integrate it fast enough to restructure (low absorption rate).

Bessen (2015) documented an instructive historical precedent that complicates the simple displacement narrative. Between 1985 and 2002, the number of automated teller machines in the United States grew from approximately 60,000 to approximately 352,000 while bank teller employment rose from approximately 485,000 to approximately 527,000. The ATM demonstrably accessed the core task of cash dispensing, yet teller employment increased because lower branch operating costs enabled banks to open more branches, and the teller role evolved toward relationship banking. The Burning Glass Institute (2025), in "The Case of the Vanishing Teller," extended this analysis through the subsequent period: teller employment has declined approximately 30 per cent since 2010, and job postings for teller positions have fallen by nearly two-thirds. The full trajectory reveals that displacement followed a different causal pathway than capability access. ATMs did not displace tellers; digital banking did, two decades later, through a mechanism (branch closure driven by mobile adoption) that was structurally independent of the original automation. The access gap and the displacement gap were governed by different variables -- precisely the independence that the Access-Displacement Framework in Section 2 formalises as the structural independence of G_A and G_D .

A harder question is whether the OAR is itself a stable constraint or one that artificial intelligence will erode. If machine learning systems eventually optimise organisational redesign, change management, and workforce reallocation, the absorption rate could accelerate in ways that current evidence cannot predict. This paper treats the OAR as empirically stable through 2026; whether it remains so as AI matures is genuinely unknown.

The objection that current stability is temporary deserves acknowledgment. OAR does not claim that displacement will never occur. It claims that displacement is rate-limited, and

that the current rate makes specific compressed timelines implausible. Klarna explicitly attempted to accelerate past the rate constraint and was forced to reverse. Enterprise artificial intelligence abandonment is currently increasing, not decreasing. Even if the rate doubles from its current baseline -- 7 per cent full scaling, approximately zero per cent institutional headcount displacement -- the resulting pace would not produce macroeconomic displacement within two years.

4.2 The Economic Forcing Function

The Organisational Absorption Rate constrains the demand side of displacement: how fast institutions can restructure around artificial intelligence. The Economic Forcing Function constrains the supply side: the cost structure of artificial intelligence deployment itself drives architectural evolution toward hybrid human-AI systems rather than centralised replacement.

The argument proceeds deductively from five empirical premises.

Premise 1: Moore's Law is dead. Dennard scaling broke down approximately 2005-2006; power density no longer decreases with transistor shrinkage. Intel required five years (2014-2019) for the 14nm to 10nm transition, not the two years Moore's Law would predict. Intel's former CEO Pat Gelsinger acknowledged in late 2023 that "we're no longer in the golden era of Moore's Law." Microprocessor performance improvement has fallen to approximately 3 per cent annually in the post-Dennard era (Hennessy and Patterson, 2019; Leiserson et al., 2020).

Premise 2: Centralised compute scaling hits an escalating cost ceiling. Without Moore's Law delivering exponential price-performance improvement, each increment of artificial intelligence capability costs more per unit than the last. Frontier model training costs have been growing at approximately 2.4 times per year since 2016 (Cottier et al., 2024). Hardware price-performance is improving at only 0.14 orders of magnitude per year -- far below the 0.5 orders of magnitude per year that the Moore's Law era delivered. The cost trajectory is not merely linear but compounding: each generation of frontier model requires disproportionately more compute for diminishing capability increments.

Premise 3: The industry has responded with miniaturisation. Every major artificial intelligence laboratory has converged on developing models with fewer than 3 billion parameters optimised for edge deployment: Meta (Llama 3.2, 1B/3B), Google (Gemma 3, 270M), Microsoft (Phi-4 mini, 3.8B), Hugging Face (SmolLM2, 135M-1.7B), Alibaba (Qwen2.5, 0.5B-1.5B), and Apple (on-device model, approximately 3B). This convergence is not one company's strategy. It is industry-wide behaviour occurring simultaneously with continued investment in frontier training -- the dual-track response to a binding cost constraint. When all participants converge on the same response to a constraint, the constraint is real and the response is economically rational.

Premise 4: Hyperscaler infrastructure has multi-decade physical life. Data centre real estate has a useful life of 20-30 years. Power infrastructure lasts 20-30 years. Cooling

systems last 15-20 years. Even if every GPU in a data centre is obsolete within three years, the shell infrastructure retains value for whatever hardware succeeds it.

Premise 5: GPUs have multi-phase useful life. Nvidia reported in 2025 that A100 GPUs shipped six years earlier were still running at full utilisation for inference workloads. The GPU lifecycle proceeds from cutting-edge training (years 1-2) to inference workloads (years 3-5) to general compute (years 5+). This is a multi-phase asset, not a single-use expenditure.

From these five premises, a conclusion follows: the economically rational architecture for artificial intelligence deployment at scale is hub-and-spoke. Centralised facilities handle heavy training. Miniaturised, quantised models deploy to edge devices for inference. Existing infrastructure becomes the hub, not a stranded asset. The cost ceiling on centralised scaling (Premises 1 and 2) forces the miniaturisation response (Premise 3), while infrastructure durability (Premises 4 and 5) ensures the hub retains value across hardware generations.

The aggregate scale of hyperscaler capital expenditure underscores the stakes. Combined spending by the five largest hyperscalers grew from approximately \$172 billion in 2022 to \$256 billion in 2024, with 2025 estimates at approximately \$427 billion and 2026 projections ranging from \$562 billion to \$700 billion (Wolf Street, February 2026; IEEE ComSoc, December 2025). Goldman Sachs projects cumulative hyperscaler capital expenditure of \$1.15 trillion from 2025 to 2027 -- more than double the \$477 billion spent from 2022 to 2024. Approximately 75 per cent of 2026 capital expenditure is allocated to artificial intelligence infrastructure (IEEE ComSoc, December 2025).

Burry (2026) argues that this infrastructure is at risk of stranding, drawing a parallel to the telecommunications fibre buildout of the late 1990s and identifying what he terms a demand-investment gap: JPMorgan analysis estimates \$650 billion in annual artificial intelligence revenue needed to deliver a 10 per cent return on modelled investments through 2030, against current generative artificial intelligence revenue of approximately \$25-37 billion -- an 18 to 26 times gap. His depreciation analysis identifies real cost pressure: hyperscalers have extended GPU depreciation schedules from 4 years to 5-6 years, collectively reducing annual depreciation charges by approximately \$18 billion. Amazon's 2025 reversal -- shortening its GPU useful life back to 5 years, taking a \$700 million operating income charge and citing "increased pace of technology development" -- constitutes a concrete admission that the 6-year assumption was aggressive.

The telecoms parallel is instructive but does not support the conclusion Burry draws. During the telecoms boom, approximately \$2 trillion was invested in 80-90 million miles of fibre optic cable. By 2001, an estimated 95 per cent was "dark" -- unused. Global telecom stocks lost over \$2 trillion in market value. The original investors were wiped out. But the infrastructure was not stranded. Level 3 Communications systematically acquired distressed fibre assets at fractions of build cost between 2003 and 2011. Verizon FiOS drove fibre-to-home demand from 2005. The iPhone (2007) required fibre backhaul for wireless networks. Netflix streaming consumed the bandwidth the fibre could provide.

Cloud computing required massive interconnection. By 2020, essentially all of the dark fibre installed during the boom was in use. Google acquired 111 Eighth Avenue in New York for \$1.8 billion in 2010 as an interconnection hub. The infrastructure that was written off in 2002 became the backbone of the modern internet -- absorbed over approximately twenty years, not stranded.

The telecoms precedent is double-edged, and intellectual honesty requires stating both edges. First, Burry's concern about equity valuations has historical support: original telecoms investors were destroyed even as the infrastructure became essential. We do not address whether hyperscaler equities are currently overvalued -- that is Burry's Thesis A, and it may have merit regardless of our analysis. Second, however, Burry's infrastructure stranding thesis -- his Thesis B -- does not follow. The infrastructure was repriced and reconsolidated, not abandoned. Artificial intelligence infrastructure has structurally better reuse potential than telecoms fibre: fibre between two points is useless if traffic does not flow between those points, whereas a data centre can serve any computational workload from any origin. The hyperscaler owners have diversified revenue and substantial free cash flow (over \$700 billion annually across the five largest), unlike the debt-laden telecoms startups of 2001. These are different claims about different questions, and conflating them weakens both.

Hennessy and Patterson (2019), in their ACM Turing Award Lecture, established the foundational observation that the end of Moore's Law creates conditions for a "New Golden Age" of architectural innovation through domain-specific hardware and hardware-software co-design. Leiserson et al. (2020) documented that post-Moore's performance gains must come from software, algorithms, and hardware specialisation, and that these gains will be "opportunistic, uneven, and sporadic." The Economic Forcing Function extends their chip-level observation to infrastructure economics: the same cost constraints that drive architectural innovation at the component level drive architectural evolution at the deployment level. The escalating cost of centralised compute creates an economic forcing function toward hybrid hub-and-spoke architectures that repurpose existing hyperscaler infrastructure as training centres while deploying miniaturised, quantised models to edge devices for inference.

This is not Schumpeterian creative destruction (Schumpeter, 1942), in which old firms are destroyed by new entrants. The Economic Forcing Function describes metamorphic evolution: the same infrastructure transforms its role from pure centralised compute to hub-and-spoke hybrid. No new entrant is required. No destruction occurs. The infrastructure metamorphoses.

Burry's data is correct. His depreciation analysis identifies real cost pressure. Cahn's (2024) "\$600 Billion Question" identified the same structural disparity between artificial intelligence infrastructure spending and revenue generation. But Burry's conclusion follows only under AGI-R economics -- the assumption, developed in Section 5, that artificial intelligence revenue derives from one-time cost savings through human replacement. Under AGI-C economics, the revenue model is recurring productivity

services: augmentation subscriptions, API access, ongoing human-AI workflow optimisation. These generate continuous returns rather than one-time savings, fundamentally altering the revenue trajectory against which the demand-investment gap is measured. The depreciation pressure Burry identifies is precisely the forcing function that drives the hybrid architecture his stranding thesis does not account for. His mechanism produces the opposite of what he predicts.

Firooz, Leduc, Liu, and Prabhakar (2025), in a Federal Reserve Bank of San Francisco Economic Letter, identified trade policy uncertainty as an additional forcing function toward automation investment. Their analysis found, however, a near-zero net employment effect from the resulting automation -- consistent with the OAR framework's prediction that even externally forced automation produces reallocation rather than net displacement at the institutional level.

Brooks' (2025) "Sin of Exponentialism" provides an honest tempering note on the forcing function's timeline. Physical deployment -- data centre construction, power grid expansion, organisational restructuring -- does not follow the exponential improvement curves of digital information processing. This supports the Economic Forcing Function's first two premises (the cost ceiling is real and physically grounded). But it also means that the hybrid pivot itself may proceed more slowly than its proponents expect. The Economic Forcing Function claims architectural direction, not timeline. The direction is established by observed industry behaviour -- universal convergence on miniaturisation, simultaneous investment in edge and frontier capabilities; the pace remains empirically open.

Two constraints, operating independently, bound the rate at which artificial intelligence can restructure labour markets. The Organisational Absorption Rate demonstrates that institutions cannot convert technological capability into workforce displacement faster than their internal complexity allows -- a finding that holds even when regulatory burden is removed, organisational size is minimised, and CEO commitment is maximised. The Economic Forcing Function demonstrates that the cost structure of computation drives deployment toward hybrid architectures that structurally embed human participation rather than eliminate it. Together, they explain why the independence of G_A and G_D documented in Section 2 persists: access gaps compress because technology diffuses rapidly, but displacement gaps do not compress because organisational and economic constraints govern a different, slower process.

If absorption constrains the rate of displacement and infrastructure economics force hybrid deployment, the question becomes: what architecture does the system evolve toward? Section 5 addresses this directly, first by naming the unstated assumption that separates the displacement and complementarian traditions (AGI-R versus AGI-C), and then by describing the deployment topology (the Hybrid Swarm Architecture) that the Economic Forcing Function produces as its architectural expression.

5. Coexistence Architecture

Section 4 identified two independent constraints on displacement speed: the Organisational Absorption Rate, which limits how fast institutions can restructure around artificial intelligence regardless of the technology's capability, and the Economic Forcing Function, which drives deployment architecture toward hybrid systems rather than centralised replacement. If absorption constrains the rate and economics force hybrid deployment, the question becomes: what does the system evolve toward? This section answers in two parts. The first names the unstated assumption that separates the displacement and complementarian traditions, transforming it from an implicit framing commitment into a testable empirical question. The second describes the deployment topology that the Economic Forcing Function produces as its architectural expression.

5.1 Redefining General Intelligence: AGI-C and AGI-R

The displacement literature surveyed in Section 1 rests on an unstated assumption about what artificial general intelligence means. When Goldman Sachs (2023) identifies 300 million jobs as "exposed" to artificial intelligence, the implicit premise is that sufficiently capable AI will substitute for human cognition across those roles. When Citrini Research (2026) project 10.2 per cent unemployment within two years, they assume a trajectory in which machine cognition renders human contribution unnecessary across broad economic domains. When Amodei (2024, 2026) characterises the human-AI "centaur phase" as "very brief," he assumes that machine capability will cross a threshold after which human partnership adds no value.

Conversely, the complementarian tradition, including Autor (2015), Bessen (2018), Brynjolfsson, Li, and Raymond (2023), and Acemoglu, Autor, and Johnson (2026), implicitly assumes that human-AI partnership will persist as the dominant mode even as machine capability increases. Neither tradition names its assumption.

We introduce two terms. **AGI-R** (artificial general intelligence as replacement) denotes the assumption that machine cognition will match or exceed human cognitive capability across all economically relevant domains, rendering human contribution unnecessary. Under AGI-R, the "general" in artificial general intelligence is a property of the machine: the machine alone achieves generality. **AGI-C** (artificial general intelligence as coexistence) denotes the alternative: that machine cognition achieves broad capability but that general intelligence is a property of the integrated human-AI system, not of either component alone. Under AGI-C, the human component provides capabilities (embodied contextual understanding, creative recombination across distant domains, normative judgment grounded in lived experience) that are not merely difficult for machines to replicate but categorically different from computational processing. This is not a semantic preference. It is a claim about the unit of analysis: whether general intelligence is assessed at the level of the machine or the level of the human-AI system. The distinction generates different predictions, different policy prescriptions, and different assessments of the empirical evidence.

Legg and Hutter (2007) catalogued over 70 definitions of artificial general intelligence. None distinguishes between replacement and coexistence as definitional commitments. Morris et al. (2023), in a Google DeepMind paper published at ICML 2024, proposed a two-dimensional framework separating performance levels from autonomy levels and explicitly stated that "increasing capabilities unlock new interaction paradigms, but do not determine them." Their framework supports the premise that capability and autonomy are independent -- but it frames collaboration as a deployment choice, not a structural property of general intelligence. Hendrycks et al. (2025) explicitly separate AGI from "Replacement AI" -- a system that renders human labour economically obsolete -- demonstrating that the serious AGI research community does not equate capability with replacement. But even Hendrycks defines AGI as a property of the machine ("an AI that can match or exceed the cognitive versatility and proficiency of a well-educated adult"). AGI-C argues that the "general" in general intelligence is a system-level property, not a component-level property. Making this distinction explicit is this paper's first contribution: it transforms an unstated framing assumption that pervades both traditions into a testable empirical question.

Five converging lines of evidence motivate AGI-C as a hypothesis.

First, biological precedent. The brain -- the only known instance of general intelligence -- achieves generality through integration of specialised subsystems: visual cortex, language centres, motor cortex, prefrontal executive function, hippocampal memory. None of these subsystems is individually general. Generality emerges from their integration. Current artificial intelligence development is producing specialised systems that parallel this architecture: large language models (language processing), vision models (image recognition), reinforcement learning agents (decision-making), robotics models (physical action). Moravec's Paradox (Moravec, 1988) -- that sensorimotor tasks trivial for humans remain difficult for artificial intelligence, while pattern-matching tasks trivial for artificial intelligence remain difficult for humans -- predicts persistent domain-dependent complementarity. The biological precedent does not prove that general intelligence requires human cognition. It establishes that the only known path to generality is subsystem integration, motivating the hypothesis that the artificial intelligence trajectory will follow the same architectural pattern rather than the monolithic scaling model.

Second, the Extended Mind Thesis. Clark and Chalmers (1998) argued that cognitive processes extend beyond the skull into external resources when those resources are reliably coupled to cognitive function. If an AI system reliably extends human cognition, and a human reliably provides contextual grounding that the AI system lacks, the combined system constitutes a single cognitive entity under the Extended Mind framework. The "general intelligence" resides in the integrated system, not in either component alone. No prior work has applied the Extended Mind Thesis to redefine artificial general intelligence specifically -- to argue that the boundary of general intelligence is functional, not anatomical, and that a reliably coupled human-AI system is a single cognitive system in the philosophically precise sense.

Third, deployment evidence. The most capable artificial intelligence deployments in 2024-2026 are overwhelmingly human-in-the-loop. McKinsey (2025) reports that only 2 per cent of organisations have deployed agentic AI at scale, with 61 per cent remaining in exploration. Gartner (2025) reports that only 15 per cent of IT application leaders are considering fully autonomous AI agents, and projects that 40 per cent of agentic AI projects will be cancelled by 2027. Deloitte (2026), surveying 3,235 respondents across 24 countries, finds that only one in five companies planning agentic artificial intelligence has mature governance frameworks in place. If AGI-R were the natural trajectory, deployment patterns would converge on autonomy as systems become more capable. They show the opposite: graduated trust calibration, in which organisations grant autonomy progressively and maintain active oversight.

Fourth, experimental evidence on human-AI interaction modes. Randazzo et al. (2025), studying 244 BCG consultants, identified three interaction patterns: Centaurs (clear division of human and AI tasks), Cyborgs (deep integration with tight feedback loops), and Self-Automators (full delegation with minimal oversight). Self-Automators (the behavioural pattern corresponding to AGI-R) developed the fewest new skills and produced the worst outcomes on both task quality and skill development. Centaurs and Cyborgs outperformed on both dimensions. Full delegation does not merely produce worse immediate outcomes; it degrades the human component of the system. Under AGI-C framing, this result is predicted: removing the human component degrades the system's general capability because the human provides contextual grounding, creative synthesis, and domain judgment that the artificial intelligence component lacks.

Fifth, the domain-dependence of human-AI advantage. Vaccaro, Almaatouq, and Malone (2024), in a meta-analysis of 106 experimental studies (370 effect sizes) published in *Nature Human Behaviour*, found that human-AI combinations performed significantly worse on average than the best of humans or AI alone -- but that this average concealed a critical domain dependence. For decision-making tasks where AI exceeds human baseline, adding humans degraded performance. For creative and generative tasks, human-AI combinations maintained their advantage. This domain dependence is precisely what AGI-C predicts: the partnership adds value on tasks requiring contextual understanding and creative synthesis, and subtracts value on narrow, well-defined tasks where computational scale is the binding factor. Naive collaboration (adding humans to everything) fails. Intelligent partnership (calibrating the human-AI balance by task domain) succeeds. The architecture of the collaboration matters, not merely its existence.

Several predecessors operate in adjacent intellectual territory and must be distinguished precisely. Dellermann et al. (2019) proposed Hybrid Intelligence: that combined human-AI systems perform better than either component alone. The distinction is between a functional claim and an ontological one. Hybrid Intelligence says combined systems *perform better*; AGI-C says general intelligence *is* the combined system. Hybrid Intelligence is pragmatic and explicitly interim ("in the next decades"); AGI-C argues partnership is the endpoint, not a transitional phase. Critically, Hybrid Intelligence cannot explain the Vaccaro finding that adding humans sometimes *degrades* performance. AGI-C

can: the human component adds value specifically on tasks requiring contextual and creative capabilities, not uniformly. The two frameworks generate different predictions, and the empirical evidence discriminates between them.

Licklider (1960) proposed man-computer symbiosis as a design goal for engineers. AGI-C argues partnership is a structural property of general intelligence, not a design aspiration. Licklider's argument is prescriptive ("build symbiotic systems"); AGI-C is descriptive ("general intelligence works this way"). Licklider expected the symbiotic phase to be temporary, giving way to machine autonomy. AGI-C argues the partnership is permanent -- not because machines will never become more capable but because generality is a system-level property that requires both components.

Wang (2025), writing in *Philosophies*, argued that human-machine hybrid systems can constitute genuine collective agents through cognitive integration -- satisfying criteria of goal alignment, functional complementarity, and stable interactivity. However, Wang addresses collective *agency* -- the capacity for joint action -- not general *intelligence*. The paper does not engage with AGI definitions. AGI-C operates at a different level of analysis: it is a claim about the nature of general intelligence, not about the capacity for coordination.

The cognitive load of centaur architecture deserves honest acknowledgment. When an artificial intelligence system performs a task in ten seconds that a human must then audit in ten minutes, the economics of partnership are not uniformly favourable. This is a real cost of AGI-C that the complementarian tradition has under-studied. It does not invalidate the framework -- the alternative (no audit, full delegation) produces the Self-Automator degradation that Randazzo et al. documented -- but it must be named as a constraint on the efficiency of human-AI systems and as a research priority for optimising the partnership architecture.

The connection to the Turing Trap strengthens the macroeconomic implications. Acemoglu and Restrepo (2019) documented the post-1987 imbalance between displacement (0.70 per cent per year) and reinstatement (0.35 per cent per year). Brynjolfsson (2022) named this dynamic the "Turing Trap": the pursuit of human-like AI that replaces workers rather than AI that creates new tasks. AGI-R is the Turing Trap: replacement that reduces wages and increases inequality. AGI-C is the escape: augmentation that creates new tasks and distributes gains more broadly. The post-1987 reinstatement slowdown documented by Acemoglu and Restrepo may reflect not a failure of technology but an over-investment in replacement-oriented development and under-investment in augmentation-oriented development. The distinction between AGI-R and AGI-C names this divergence at the level of the intelligence concept itself.

The strongest counter-argument comes from Amodei (2026), who stated on the Interesting Times with Ross Douhat podcast (New York Times, 12 February 2026) that the current "centaur phase" will be "very brief," and separately estimated that artificial intelligence could replace 50 per cent of entry-level white-collar tasks within two to three years (Axios, January 2026). His argument draws on the chess precedent, in which engines surpassed

centaur teams within approximately a decade. Three responses are warranted. First, chess is a closed, finite-state domain with perfect information and deterministic rules -- approximately 10^{47} legal positions within a bounded state space; general intelligence operates in open, unbounded domains with incomplete information, contextual ambiguity, and normative judgment. The tractability of chess to pure computation does not generalise to open-domain cognition any more than the history of calculator development predicted the history of mathematics. Amodei provides no mechanism for this transfer, no account of how capabilities demonstrated in closed domains extend to open ones, and no timeline with specific, falsifiable benchmarks. Second, the meta-analytic evidence (Vaccaro et al., 2024) demonstrates domain dependence, not universal erosion: the centaur advantage persists in creative and generative tasks even as it erodes in narrow decision tasks. Third, the commercial context warrants acknowledgment: Anthropic's business model is oriented toward building progressively more autonomous systems, and predictions favouring autonomy over partnership align with that orientation. This observation does not invalidate the argument but is relevant context for evaluating timeline claims made without falsifiable benchmarks.

AGI-C is falsifiable. It fails if fully autonomous systems consistently outperform human-AI systems across open-ended, creative, contextually-embedded, and normatively-loaded tasks; if enterprise deployment evidence shifts decisively toward autonomous outperformance rather than graduated trust calibration; or if the miniaturisation trend documented in Section 4.2 reverses in favour of exclusively centralised deployment. As of 2026, the empirical trajectory supports AGI-C across all three dimensions.

5.2 Hybrid Swarm Architecture

AGI-C identifies the conceptual trajectory: coexistence rather than replacement. This section identifies its architectural expression. The Economic Forcing Function (Section 4.2) established that the death of Moore's Law creates an escalating cost ceiling on centralised computation, and that the industry has responded with miniaturisation. The question is what deployment topology this response produces. The answer is already observable: a three-layer hybrid system that is not one option among many but the economically forced outcome of the constraints Section 4.2 documented.

The architecture comprises three layers. The **centralised layer** handles heavy training, frontier research, and complex multi-step reasoning -- tasks requiring GPU clusters and large parameter models. The **distributed layer** handles edge inference, real-time adaptation, and localised intelligence -- tasks running on quantised, miniaturised models deployed to billions of edge devices. The **integration layer** operates bidirectionally: model compression and knowledge distillation (Hinton, Vinyals, and Dean, 2015) flow downward from central training to edge deployment; federated learning (McMahan et al., 2017) and aggregated updates flow upward from edge experience to central model improvement. The causal chain from Section 4.2 produces this architecture deductively: the cost ceiling (Premises 1 and 2) forces the miniaturisation response (Premise 3), which populates the

distributed layer; infrastructure durability (Premises 4 and 5) ensures the centralised hub retains value across hardware generations; the integration layer connects them.

Apple Intelligence provides the first mass-market production implementation. Apple's on-device model -- approximately 3 billion parameters, trained with 2-bit quantisation-aware training and optimised for Apple Silicon -- handles privacy-sensitive and latency-critical tasks locally. Private Cloud Compute (Apple, 2024) processes requests exceeding on-device capability, with cryptographic guarantees that user data is neither stored nor accessible. A routing layer dynamically allocates requests between device and cloud based on complexity and privacy requirements. The system is deployed to approximately 2 billion active Apple devices worldwide. This is not a research prototype. It is a production system implementing all three layers at planetary scale, arrived at for precisely the economic and privacy reasons the framework predicts: on-device inference is cheaper per query, faster (no network latency), and more private; cloud compute handles what devices cannot.

Google's federated learning deployments demonstrate the integration layer in production. Gboard's next-word prediction is trained across millions of Android devices, with raw text never leaving the device. Google achieved the first production neural network with a formal differential privacy guarantee -- the DP-FTRL algorithm deployed for Spanish-language Gboard -- demonstrating that privacy-preserving bidirectional integration is productionised at scale, not merely theoretically viable.

The miniaturisation convergence documented in Section 4.2 supplies the distributed layer with its material substrate. Every major artificial intelligence laboratory has simultaneously invested in sub-3 billion parameter models optimised for edge deployment: Meta (Llama 3.2, 1B/3B), Google (Gemma 3, 270M), Microsoft (Phi-4 mini, 3.8B), Hugging Face (SmolLM2, 135M-1.7B), Alibaba (Qwen2.5, 0.5B-1.5B). Compression techniques have advanced to 2-bit quantisation-aware training, 1.58-bit ternary quantisation, and hybrid pruning-quantisation-distillation pipelines achieving 75 per cent model size reduction with 97 per cent accuracy retention. Phi-4 at 14 billion parameters achieves 84.8 per cent on the MATH benchmark while running locally at 15 times the speed of frontier models. The market data reflects the trajectory: the edge AI market reached \$24.91 billion in 2025 with a projected compound annual growth rate of 21.7 per cent to \$118.69 billion by 2033. Inference workloads consume over 55 per cent of AI-optimised infrastructure spending in early 2026, projected to reach 70-80 per cent by year-end. The federated learning market, at \$0.1 billion in 2025, is projected to reach \$1.6 billion by 2035 -- small in absolute terms but growing at 27.3 per cent compound annually, reflecting the nascent but accelerating adoption of the integration layer beyond its Google and Apple pioneers.

Two prior treatments of distributed artificial intelligence require engagement. Gharbawi (2025), writing on the Bank of England's Bank Underground blog, proposed that artificial general intelligence may emerge from "a constellation of interacting AI agents" through bottom-up complex systems dynamics, drawing on biological analogies of ant colonies and neural networks. Gharbawi does not specify an architecture, does not address

compute economics or Moore's Law, does not discuss model compression or distillation, and provides no empirical evidence beyond biological analogies and existing multi-agent demonstrations. The Hybrid Swarm Architecture differs on every substantive dimension: it proposes a specific three-layer hub-and-spoke topology, identifies the Economic Forcing Function as the causal mechanism, grounds its claims in production deployment data and hyperscaler capital expenditure figures, and connects to the infrastructure investment analysis that Gharbawi's conceptual treatment omits entirely.

Tomasev et al. (2025), from Google DeepMind, proposed the "patchwork AGI hypothesis": that AGI-level capability may emerge from coordination of sub-AGI agents through market mechanisms, and identified six categories of emergent safety risk. Their paper addresses a fundamentally different question. Tomasev et al. ask what safety implications arise if AGI emerges spontaneously from distributed agents. The Hybrid Swarm Architecture asks what deployment topology compute economics forces. The fundamental distinction is between emergence and engineering. DeepMind's patchwork is emergent: bottom-up, unplanned, agents self-organising through market dynamics. The Hybrid Swarm is engineered: top-down, designed, driven by cost ceilings that leave no economically rational alternative. DeepMind addresses what happens when agents coordinate spontaneously; the Hybrid Swarm describes what happens when engineers deliberately design hub-and-spoke deployment because the economics demand it.

The emergence question itself deserves honest treatment. No engineered distributed system has demonstrated emergent general intelligence. Multi-agent coordination produces task-level capability gains, but these are domain-specific optimisations, not emergent generality. The Hybrid Swarm Architecture stands as an architectural and economic thesis independent of whether the distributed topology produces emergent capabilities beyond the sum of its parts. This is the framework's weakest point, and acknowledging it is necessary. The architectural thesis holds regardless of the emergence question: even without emergence, the three-layer topology is more efficient, more scalable, and more economically rational than purely centralised deployment. The emergence question remains empirically open, and we note that leading artificial intelligence safety researchers consider the possibility sufficiently plausible to warrant proactive governance (Tomasev et al., 2025).

The causal chain connecting Sections 4 and 5 is now complete. The death of Moore's Law creates an escalating cost ceiling on centralised computation (Section 4.2, Premises 1-2). The industry responds rationally with miniaturisation (Premise 3), populating the distributed layer. Infrastructure durability (Premises 4-5) ensures the centralised hub retains value. The integration layer connects them. The Hybrid Swarm Architecture is the economically forced expression of the cost dynamics that the Economic Forcing Function identifies, and the architectural embodiment of the coexistence that AGI-C describes. Under this architecture, existing hyperscaler infrastructure is not stranded; it is the hub. The Burry infrastructure stranding thesis, introduced in Section 1 and addressed empirically in Section 4.2, finds its architectural resolution: the depreciation pressure Burry identifies is precisely the forcing function that drives the architectural evolution his

stranding thesis does not account for. His mechanism produces the opposite of what he predicts.

Together, AGI-C and the Hybrid Swarm Architecture provide the positive case that the preceding sections' negative arguments -- measurement is degrading, absorption is rate-limited, economics force hybrid deployment. These negative arguments imply a positive case but do not state it directly. The displacement literature assumes AGI-R. The empirical trajectory points to AGI-C. The infrastructure being built implements hybrid coexistence, not centralised replacement. If this analysis is correct, the question for policymakers and institutions is not how to prepare for mass displacement but how to invest in the human-AI partnership that the economic and architectural evidence indicates is the dominant trajectory. Section 6 addresses the policy implications, cross-domain implications of skill compression, the limitations of this framework, and directions for future research.

6. Discussion

The preceding sections have constructed a layered argument. Section 2 demonstrated that capability access and economic displacement are structurally independent, with cross-generational evidence showing that the access gap has compressed from decades to years while displacement gaps have not followed. Section 3 documented two mechanisms through which the measurement infrastructure underpinning displacement predictions is actively degrading: the Measurement Obsolescence Hypothesis (establishment surveys overestimating traditional employment) and the Sampling Residualisation Hypothesis (household surveys underestimating adaptive employment). Section 4 identified two rate-limiting mechanisms, the Organisational Absorption Rate and the Economic Forcing Function, that constrain displacement speed regardless of measurement accuracy. Section 5 named the unstated assumption dividing the displacement and complementarian traditions (AGI-R versus AGI-C) and described the Hybrid Swarm Architecture as the economically forced deployment topology. This section draws out the implications, extends the framework's most robust empirical finding to an unstudied domain, acknowledges the framework's limitations honestly, and identifies the empirical programme required to test it.

6.1 Policy Implications

Displacement predictions are not merely academic exercises. They drive policy. Frey and Osborne's (2013) estimate that 47 per cent of United States occupations faced high automation risk informed a generation of workforce retraining programmes, universal basic income proposals, and automation taxation debates. Goldman Sachs's (2023) identification of 300 million "exposed" jobs shaped institutional risk assessments. Citrini Research's (2026) two-year timeline, however questionable its methodology, contributed to intraday equity volatility. When policymakers accept displacement predictions calibrated to access-gap dynamics, they allocate resources to the wrong problem on the wrong timescale.

If the Access-Displacement Framework holds, three policy recalibrations follow.

First, the rate-limiting factor in labour market adjustment is organisational, not technological. Section 4 demonstrated that even JPMorgan Chase -- the institution with the largest technology budget in global banking, maximum CEO commitment, and explicit willingness to restructure -- produced redeployment rather than displacement. Klarna, operating without JPMorgan's regulatory constraints, attempted to accelerate past the absorption rate and was forced to reverse. Policy designed to slow technological adoption addresses a constraint that is not binding. Policy designed to accelerate organisational absorption, by reducing the friction of internal redeployment, supporting complementary innovation, and investing in management capacity for technology integration, addresses the actual bottleneck. The distinction is between building a dam and widening a channel.

Second, the measurement crisis documented in Section 3 means that policymakers are navigating with degrading instruments. Bureau of Labor Statistics benchmark revisions have escalated from historical norms to 3.4 times the historical average in the absence of any declared recession. The Current Population Survey response rate has reached a series low of 64 per cent. If the employment-to-population ratio is five percentage points higher than measured, as Goldman Sachs estimates, then monetary policy calibrated to understated employment data risks systematic miscalibration. The policy implication is not that measurement should be abandoned but that it must be supplemented with the Novel Temporal Labour Indicators that Section 3 motivated, including instruments like Soto's (2025) AI Research Index that capture dynamics establishment surveys cannot detect.

Third, the distinction between AGI-R and AGI-C has direct policy consequences. If the dominant trajectory is replacement, the appropriate policy response is compensatory: income transfers, retraining, safety nets. If the dominant trajectory is coexistence, the appropriate response is generative: investment in the human-AI partnership infrastructure that Brynjolfsson (2022) identified as under-funded, the new task creation that Acemoglu and Restrepo (2019) documented as having slowed, and the institutional frameworks that translate augmentation into broadly distributed productivity gains. Acemoglu, Autor, and Johnson (2026), in the most comprehensive recent treatment, proposed a five-category taxonomy of technological change and identified three market failures biasing artificial intelligence development toward automation rather than augmentation: insufficient consideration of the task-level composition of automation decisions, inadequate investment in new task creation, and information failures that cause firms to over-automate. Their "pro-worker AI" framework, distinguishing labour-augmenting, capital-augmenting, automating, expertise-levelling, and new-task-creating innovations, provides the policy architecture that the AGI-C trajectory requires. Compensatory policy presumes AGI-R; generative policy presumes AGI-C. The empirical evidence presented in this paper supports AGI-C, and policy should be designed accordingly -- not as a prediction but as the trajectory the evidence currently supports, and one that policy can actively reinforce.

The framework does not prescribe specific policy instruments. It identifies what the evidence implies for the *type* of policy response. The consistent finding across Sections 2 through 5 is that the binding constraints on artificial intelligence's labour market impact

are institutional and economic, not technological. Policy that addresses technological capability without addressing institutional capacity misidentifies the problem.

6.2 Cross-Domain Skill Compression

The distributional compression documented in the experimental literature -- lower performers gaining the most from artificial intelligence augmentation, higher performers gaining the least -- is the single most robust empirical finding across all experimental studies of artificial intelligence in the workplace. Brynjolfsson, Li, and Raymond (2023) found that the lowest-skilled customer support agents gained approximately 35 per cent in productivity, while the highest-skilled agents gained negligibly. Dell'Acqua et al. (2023) documented a 43 per cent quality improvement for below-average management consultants versus 17 per cent for above-average consultants. Noy and Zhang (2023) found similar distributional effects for professional writing tasks. The pattern is consistent: artificial intelligence compresses the skill distribution within a domain by raising the floor substantially while leaving the ceiling largely unchanged.

Every study documenting this effect measures it *within* a single professional domain. A customer support agent augmented by artificial intelligence becomes a more productive customer support agent. A management consultant augmented by artificial intelligence becomes a more productive management consultant. The skill ceiling against which compression is measured is domain-specific: years of experience, tacit knowledge, and pattern recognition within that particular field.

A hypothesis follows that existing research has not investigated. Consider an individual who is a high performer in Domain A -- a financial markets practitioner with two decades of institutional experience, for instance -- who applies artificial intelligence augmentation in Domain B, such as academic research and writing. Within Domain A, this individual operates near the skill ceiling; artificial intelligence augmentation provides marginal gains, consistent with the experimental findings. Within Domain B, the same individual is not at ceiling. They possess the intellectual capacity, analytical rigour, and pattern recognition that made them a high performer in Domain A, but they lack the domain-specific training -- the tacit knowledge of academic conventions, the familiarity with literature across multiple disciplines, the facility with formal academic articulation -- that defines high performance in Domain B. In Domain B, this individual is mid-range at best. And it is precisely in the mid-range where artificial intelligence augmentation compresses most.

This hypothesis carries an uncomfortable corollary. If artificial intelligence enables any competent professional to perform at mid-range level in domains beyond their training, the wage premium for domain-specific expertise may erode even as total employment holds steady. Cross-domain skill compression could preserve jobs while undermining the economic returns to specialisation. Whether this constitutes augmentation or a subtler form of displacement is a question the framework does not resolve.

The implication is that the skill ceiling is *domain-specific, not person-specific*. A high performer crossing domains becomes a mid-range performer in the new domain -- and

mid-range performers gain the most from augmentation. Cross-domain transfer for experts may be artificial intelligence's largest value proposition for high performers, precisely because existing studies do not measure it. The existing literature measures within-domain compression, where high performers gain little. The unstudied phenomenon is cross-domain compression, where high performers gain substantially because they are no longer at ceiling.

I present this as a hypothesis requiring formal investigation, not as a finding. A controlled study comparing within-domain augmentation gains for high performers against cross-domain augmentation gains for the same individuals would test the claim directly. If confirmed, it would extend the most robust finding in the artificial intelligence augmentation literature to its least-studied and potentially most consequential dimension.

6.3 Limitations

This paper advances seven contributions as testable hypotheses. I acknowledge openly what the evidence cannot yet establish.

The cross-generational evidence supporting the Access-Displacement Framework (Section 2) comprises five technology transitions. Statistical independence cannot be established from five observations. The framework is presented as a hypothesis with illustrative evidence, not as a proven statistical relationship. Comin and Hobijn's (2004; 2010) cross-country analysis of 104 technologies provides supporting evidence from a substantially larger dataset, but their methodology addresses adoption speed convergence, not the specific G_A/G_D independence claim.

The Measurement Obsolescence Hypothesis rests on two consecutive outsized benchmark revisions (March 2024 and March 2025) against a baseline of 22 annual observations. Two anomalous data points, however directionally consistent, do not constitute proof of structural obsolescence. The falsification condition -- the March 2026 benchmark revision -- provides the confirmatory test but will not be available until approximately August 2026.

The Sampling Residualisation Hypothesis identifies a plausible mechanism for systematic survey nonresponse bias but relies on indirect evidence. The core empirical test -- regressing state-level CPS nonresponse rates against independent work growth -- requires data that is available within Census Bureau systems but not publicly accessible. Until that test is conducted, the SRH remains a hypothesis with supporting circumstantial evidence.

AGI-C specifies falsification conditions but acknowledges that falsification in real time is difficult. The claim that general intelligence is a system-level property rather than a component-level property may resist empirical disconfirmation for years, as autonomous systems would need to demonstrate consistent outperformance across open-ended, creative, contextually embedded, and normatively loaded tasks -- domains where benchmarking is itself underdeveloped.

The Organisational Absorption Rate argument rests on an upper bound derived from a single institution (JPMorgan Chase), corroborated by Klarna's reversal and enterprise survey data. The upper bound logic is sound -- one data point suffices for a ceiling argument -- but the generalisability of the specific rate constraint across sectors, firm sizes, and regulatory environments remains empirically open.

The Hybrid Swarm Architecture stands as an architectural and economic thesis independent of whether the distributed topology produces emergent capabilities beyond the sum of its parts. No engineered distributed system has demonstrated emergent general intelligence. The architectural claim holds without emergence; the stronger claim that the architecture will produce emergent general capability remains speculative.

The empirical foundation of this paper is predominantly United States-centric. BLS benchmark data, CPS nonresponse rates, and the majority of enterprise case studies are drawn from the American labour market. Whether the Access-Displacement Framework's independence claim, the MOH's measurement degradation pattern, and the OAR's rate constraint generalise to economies with different labour market structures -- the United Kingdom, the European Union, East Asia, emerging markets -- is untested.

Enterprise artificial intelligence adoption data cited throughout is 2024-2025 vintage. The landscape is evolving rapidly. Findings that 88 per cent of organisations have adopted artificial intelligence but only 7 per cent have fully scaled (McKinsey, 2025) may look different by 2027.

There is also a political economy dimension the paper does not address. If the framework is correct that displacement predictions are systematically overstated, the policy infrastructure built around those predictions, from retraining programmes to universal basic income pilots, may itself become a vested interest resistant to updating. Institutions designed for an emergency that does not materialise rarely disband quietly.

The framework focuses on aggregate employment dynamics and does not address the quality of redeployment. If "redemption" means transition from stable employment with benefits to gig work, independent contracting, or lower-compensated positions, the framework's aggregate optimism requires substantial qualification. The distinction between employment *quantity* and employment *quality* is a limitation that future work must address directly.

Unresolved liability frameworks for autonomous artificial intelligence decisions constitute an additional absorption constraint that reinforces the Organisational Absorption Rate mechanism but is not developed in this paper.

Finally, this is a hypothesis paper, not an empirical paper. I propose testable frameworks and identify the evidence that would confirm or falsify them. The paper does not conduct the definitive tests.

6.4 Future Work

The framework generates a specific programme of empirical investigation across multiple timescales.

Near-term (2026-2027). The March 2026 benchmark revision, expected in preliminary form by approximately August 2026, provides the first confirmatory or falsifying test for the Measurement Obsolescence Hypothesis. If the revision returns to within one standard deviation of the 2002-2023 historical mean, the MOH is substantially weakened. If it continues the escalating pattern, the MOH gains significant predictive credibility.

State-level CPS nonresponse data, obtainable through Census Bureau collaboration or Freedom of Information Act request, would enable the core test of the Sampling Residualisation Hypothesis: whether states with the fastest growth in independent work show the largest increases in survey nonresponse.

Industry-level benchmark analysis -- comparing CES revision patterns in AI-exposed sectors (professional services, information technology, finance) against less-exposed sectors (construction, mining, agriculture) -- would test whether measurement degradation is concentrated where artificial intelligence adoption is highest.

Medium-term (2026-2030). Longitudinal OAR tracking across institutions from 2022 to 2028 would test whether the organisational absorption rate is accelerating, decelerating, or stable as artificial intelligence tools mature. The prediction is stability or deceleration: firms investing more in artificial intelligence but not displacing proportionally more workers.

Cross-country comparison of employment survey benchmark revisions -- comparing BLS revisions to United Kingdom ONS Labour Force Survey corrections, Australian Bureau of Statistics revisions, Statistics Canada benchmarks, and Eurostat equivalents -- would test whether the measurement degradation pattern is specific to the United States or correlated with artificial intelligence adoption rates internationally.

A critical question is whether the structural independence of G_A and G_D holds at higher adoption levels. Current evidence covers a 20 to 40 per cent adoption range (Crane, Green, and Soto, 2025). If artificial intelligence adoption exceeds 40 per cent of the workforce, the institutional absorption mechanisms that maintain G_A/G_D independence may be overwhelmed. This threshold effect is explicitly testable with longitudinal data as adoption increases and constitutes the most important boundary condition on the framework's central claim.

Formal investigation of cross-domain skill compression, as outlined in Section 6.2, would extend the most robust finding in the artificial intelligence augmentation literature to its least-studied dimension. A controlled study comparing within-domain and cross-domain augmentation gains for matched high-performing individuals across multiple professional domains would provide direct evidence.

Long-term. A systematic AGI-C empirical programme -- comparing centaur versus fully autonomous performance across open-ended, creative, contextually embedded, and normatively loaded tasks, tracked longitudinally as artificial intelligence capability increases -- would test whether the domain-dependent centaur advantage identified by Vaccaro et al. (2024) persists, erodes, or strengthens over time. This is the definitive test of the AGI-C hypothesis and will require years of accumulated data.

A second paper in the publication sequence develops the full philosophical treatment of AGI-C that this paper introduces: the extended cognition literature, the phenomenology of human-AI systems, consciousness and intentionality in integrated systems, and the implications of AGI-C for artificial intelligence safety. A third paper addresses an independent empirical contribution in a related domain, extending the theoretical foundations established here.

Acknowledgments

The author acknowledges the use of artificial intelligence tools (large language models) in the preparation of this paper. These tools were used for literature search and synthesis across multiple academic disciplines, citation verification and cross-referencing, copy-editing, and formatting. All theoretical frameworks, hypotheses, original contributions, and analytical arguments are the author's own work, developed from two decades of institutional market experience and cross-disciplinary observation. The author reviewed and edited all AI-assisted output and takes full responsibility for the content of this paper.

References

- Abraham, K.G., Haltiwanger, J.C., Hou, C., Sandusky, K. and Spletzer, J.R. (2021). Reconciling survey and administrative measures of self-employment. *Journal of Labor Economics*, 39(4), 825-860.
- Acemoglu, D. (2024). The simple macroeconomics of AI. NBER Working Paper No. 32487.
- Acemoglu, D. and Restrepo, P. (2019). Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives*, 33(2), 3-30.
- Acemoglu, D., Autor, D. and Johnson, S. (2026). Building pro-worker artificial intelligence. NBER Working Paper No. 34854. DOI: 10.3386/w34854.
- ADP Research Institute (2025). Workforce trends data: Payroll job creation slowdown and wage growth analysis. ADP.
- Allen, R.C. (2009). Engels' pause: Technical change, capital accumulation, and inequality in the British industrial revolution. *Explorations in Economic History*, 46(4), 418-435.
- Amodei, D. (2024). Machines of loving grace. darioamodei.com.

Amodei, D. (2026). Remarks on the centaur phase. Interesting Times with Ross Douthat podcast (New York Times), 12 February 2026. Entry-level task estimate from Axios interview, January 2026.

Apple (2024). Private Cloud Compute. apple.com/privacy/docs/Apple_Private_Cloud_Compute.pdf.

Arntz, M., Gregory, T. and Zierahn, U. (2016). The risk of automation for jobs in OECD countries. OECD Social, Employment and Migration Working Papers No. 189.

Autor, D.H. (2015). Why are there still so many jobs? The history and future of workplace automation. *Journal of Economic Perspectives*, 29(3), 3-30.

Autor, D.H., Levy, F. and Murnane, R.J. (2003). The skill content of recent technological change: An empirical exploration. *The Quarterly Journal of Economics*, 118(4), 1279-1333.

Autor, D.H. and Thompson, N. (2025). Expertise. *Journal of the European Economic Association*, 23(4), 1203-1271.

Bernhardt, R., Munro, D. and Wolcott, E.L. (2024). How does the dramatic rise of nonresponse in the Current Population Survey impact labor market indicators? *Journal of Applied Econometrics*, 39(3), 498-512.

Bessen, J.E. (2015). *Learning by Doing: The Real Connection Between Innovation, Wages, and Wealth*. Yale University Press.

Bessen, J.E. (2018). AI and jobs: The role of demand. NBER Working Paper No. 24235.

Bloomberg (2025). Klarna turns from AI to real person customer service. 8 May 2025.

Bracha, A. and Burke, M.A. (2023). Informal work and official employment statistics: What's missing? Federal Reserve Bank of Boston Working Paper 23-15.

Bresnahan, T.F. and Trajtenberg, M. (1995). General purpose technologies: 'Engines of growth'? *Journal of Econometrics*, 65(1), 83-108.

Brooks, R. (2025). Predictions scorecard, 2025 January 01. rodneybrooks.com.

Brynjolfsson, E. (2022). The Turing Trap: The promise and peril of human-like artificial intelligence. *Daedalus*, 151(2), 272-287.

Brynjolfsson, E., Li, D. and Raymond, L. (2023). Generative AI at work. NBER Working Paper No. 31161.

Brynjolfsson, E., Rock, D. and Syverson, C. (2021). The productivity J-curve: How intangibles complement general purpose technologies. *American Economic Journal: Macroeconomics*, 13(1), 333-372.

Bureau of Labor Statistics. CES Birth-Death Model Historical Data. bls.gov/web/empisit/cesbdhst.htm.

Bureau of Labor Statistics. CES National Benchmark Articles (annual). bls.gov/web/empisit/cesbmart.htm.

Bureau of Labor Statistics. CPS Response Rate documentation. bls.gov/cps/methods/response_rates.htm.

Burning Glass Institute (2025). The case of the vanishing teller: How banking's entry level jobs are transforming.

Burry, M. (2026). X posts and 'Cassandra Unchained' (Substack). AI infrastructure depreciation and short positions: NVDA puts, PLTR puts.

Butcher, S. (2026). JPMorgan's AI coding efficiencies. eFinancialCareers, 23 February 2026.

Cahn, D. (2024). AI's \$600B question. Sequoia Capital, 20 June 2024.

Caselli, F. (1999). Technological revolutions. *American Economic Review*, 89(1), 78-102.

Census Bureau (2025). Using administrative data to evaluate nonresponse bias in the 2025 Current Population Survey Annual Social and Economic Supplement. Research Matters Blog, 9 September 2025.

Center on Budget and Policy Priorities. Funding for key federal statistical agencies has fallen substantially since 2010. cbpp.org.

Challapally, A., Pease, C., Raskar, R. and Chari, P. (2025). The GenAI divide: State of AI in business 2025. MIT Project NANDA, July 2025.

Citrini Research / van Geelen, J. and Shah, A. (2026). The 2028 global intelligence crisis. citriniresearch.com.

Clark, A. and Chalmers, D.J. (1998). The extended mind. *Analysis*, 58(1), 7-19.

Cohen, W.M. and Levinthal, D.A. (1990). Absorptive capacity: A new perspective on learning and innovation. *Administrative Science Quarterly*, 35(1), 128-152.

Comin, D. and Hobijn, B. (2004). Cross-country technology adoption: Making the theories face the facts. *Journal of Monetary Economics*, 51(1), 39-83.

Comin, D. and Hobijn, B. (2010). An exploration of technology diffusion. *American Economic Review*, 100(5), 2031-2059.

Congressional Research Service (2025). Current Employment Survey benchmark revisions. Report IF12827.

Cottier, B., Rahman, R., Fattorini, L., Maslej, N., Besiroglu, T. and Owen, D. (2024). The rising costs of training frontier AI models. arXiv:2405.21015.

Crane, L., Green, M. and Soto, P.E. (2025). Measuring AI uptake in the workplace. *FEDS Notes*, Board of Governors of the Federal Reserve System, 5 February 2025. DOI: 10.17016/2380-7172.3724.

David, P.A. (1990). The dynamo and the computer: An historical perspective on the modern productivity paradox. *American Economic Review*, 80(2), 355-361.

Dell'Acqua, F., McFowland, E. III, Mollick, E.R., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., Kraymer, L., Candelon, F. and Lakhani, K.R. (2023). Navigating the jagged technological frontier: Field experimental evidence of the effects of AI on knowledge worker productivity and quality. Harvard Business School Working Paper No. 24-013.

Dellermann, D., Ebel, P., Soellner, M. and Leimeister, J.M. (2019). Hybrid intelligence. *Business & Information Systems Engineering*, 61(5), 637-643.

Deloitte (2026). The state of AI in the enterprise: The untapped edge. Deloitte AI Institute. Survey: August-September 2025, 3,235 respondents, 24 countries, 6 industries.

Dimon, J. (2025). Bloomberg TV interview, 7 October 2025. JPMorgan AI investment and return data.

Fichman, R.G. and Kemerer, C.F. (1999). The illusory diffusion of innovation: An examination of assimilation gaps. *Information Systems Research*, 10(3), 255-275.

Firooz, H., Leduc, S., Liu, Z. and Prabhakar, D.B. (2025). Will trade uncertainty boost automation? *FRBSF Economic Letter* 2025-20, Federal Reserve Bank of San Francisco, 2 September 2025.

Forrester Research (2026). The Forrester AI Job Impact Forecast, US, 2025-2030 (Report No. RES190071), 13 January 2026. See also Summers, B. (2025). Predictions 2026: The Future of Work (Report No. RES185020). Forrester Research.

Frey, C.B. (2019). *The Technology Trap: Capital, Labor, and Power in the Age of Automation*. Princeton University Press.

Frey, C.B. and Osborne, M.A. (2013). The future of employment: How susceptible are jobs to computerisation? Oxford Martin School Working Paper. Published as: Frey, C.B. and Osborne, M.A. (2017). *Technological Forecasting and Social Change*, 114, 254-280.

Gartner (2025). IT application leaders survey on autonomous AI agents.

Gharbawi, M. (2025). The gathering swarm: Emergent AGI and the rise of distributed intelligence. Bank Underground (Bank of England staff blog), 21 August 2025.

Goldman Sachs / Briggs, J. and Kodnani, D. (2023). The potentially large effects of artificial intelligence on economic growth. Economics Research.

Goldman Sachs / Rindels, J. (2025). The gig economy: Another perspective on the labor market. Global Investment Research, 17 November 2025.

Greenwood, J. and Yorukoglu, M. (1997). 1974. *Carnegie-Rochester Conference Series on Public Policy*, 46, 49-95.

Hendrycks, D., Song, D., Szegedy, C., Lee, H. et al. (2025). A definition of AGI. arXiv:2510.18212.

Hennessy, J.L. and Patterson, D.A. (2019). A new golden age for computer architecture. *Communications of the ACM*, 62(2), 48-60.

Hinton, G., Vinyals, O. and Dean, J. (2015). Distilling the knowledge in neural networks. arXiv:1503.02531.

IEEE ComSoc Technology Blog (2025). Hyperscaler capex and AI allocation data. 22 December 2025.

JPMorgan Chase & Co. (2026). Company Update 2026. 23 February 2026.

Keynes, J.M. (1930). Economic possibilities for our grandchildren. In *Essays in Persuasion*. Macmillan.

Lane, P.J., Koka, B.R. and Pathak, S. (2006). The reification of absorptive capacity: A critical review and rejuvenation of the construct. *Academy of Management Review*, 31(4), 833-863.

Leduc, S., Oliveira, L.E. and Paulson, C.M. (2025). Do low survey response rates threaten data dependence? *FRBSF Economic Letter* 2025-07, 31 March 2025.

Legg, S. and Hutter, M. (2007). A collection of definitions of intelligence. *Frontiers in Artificial Intelligence and Applications*, 157, 17-24.

Leiserson, C.E., Thompson, N.C., Emer, J.S., Kuzmaul, B.C., Lampson, B.W., Sanchez, D. and Schardl, T.B. (2020). There's plenty of room at the top: What will drive computer performance after Moore's Law? *Science*, 368(6495), eaam9744.

Leontief, W. (1983). The long-term impact of technology on employment and unemployment. *National Academy of Engineering Symposium*.

Licklider, J.C.R. (1960). Man-computer symbiosis. *IRE Transactions on Human Factors in Electronics*, HFE-1(1), 4-11.

MBO Partners (2025). State of Independence in America. Annual report.

McKinsey & Company (2022). Freelance, side hustles, and gigs: Many more Americans have become independent workers. Survey of 25,062 adults by Ipsos.

McKinsey & Company (2025). The state of AI in 2025: Global survey. March 2025.

McMahan, H.B., Moore, E., Ramage, D., Hampson, S. and Arcas, B.A.y. (2017). Communication-efficient learning of deep networks from decentralized data. Proceedings of AISTATS 2017.

Moravec, H. (1988). *Mind Children: The Future of Robot and Human Intelligence*. Harvard University Press.

Morris, M.R., Sohl-Dickstein, J., Fiedel, N., Warkentin, T., Dafoe, A., Faust, A., Farabet, C. and Legg, S. (2023). Levels of AGI for operationalizing progress on the path to AGI. arXiv:2311.02462. Published at ICML 2024.

Noy, S. and Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654), 187-192.

Nvidia (2025). A100 GPU lifecycle and inference utilisation data. Nvidia corporate disclosures, 2025.

OECD (2023). *OECD Employment Outlook 2023: Artificial Intelligence and the Labour Market*. OECD Publishing.

Peng, S., Kalliamvakou, E., Cihon, P. and Demirer, M. (2023). The impact of AI on developer productivity: Evidence from GitHub Copilot. arXiv:2302.06590.

Perez, C. (2002). *Technological Revolutions and Financial Capital: The Dynamics of Bubbles and Golden Ages*. Edward Elgar.

PwC (2018). Will robots really steal our jobs? An international analysis of the potential long term impact of automation. PricewaterhouseCoopers.

Randazzo, E., Lifshitz-Assaf, H., Kellogg, K., Dell'Acqua, F., Mollick, E.R., Candelon, F. and Lakhani, K.R. (2025). Cyborgs, centaurs and self-automators: The three modes of human-GenAI knowledge work. Harvard Business School Working Paper No. 26-036.

Rodrik, D. (2016). Premature deindustrialization. *Journal of Economic Growth*, 21(1), 1-33.

Rubin, D.B. (1976). Inference and missing data. *Biometrika*, 63(3), 581-592.

S&P Global (2025). Enterprise AI adoption and abandonment survey.

Schumpeter, J.A. (1942). *Capitalism, Socialism and Democracy*. Harper & Brothers.

Soto, P.E. (2025). Research in commotion: Measuring AI research and development through conference call transcripts. *Finance and Economics Discussion Series 2025-011*, Board of Governors of the Federal Reserve System.

Tomasev, N., Franklin, M., Jacobs, J., Krier, S. and Osindero, S. (2025). Distributional AGI safety. arXiv:2512.16856.

Vaccaro, M., Almaatouq, A. and Malone, T.W. (2024). When combinations of humans and AI are useful: A systematic review and meta-analysis. *Nature Human Behaviour*, 8(12), 2293-2303.

Wang, R. (2025). Cognitive integration for hybrid collective agency. *Philosophies*, 10(5), 103. DOI: 10.3390/philosophies10050103.

Ward, J.M. and Edwards, K.A. (2021). CPS nonresponse during COVID-19. *Labour Economics*, 72, 102060.

Wells Fargo / Scharf, C.W. (2026). CEO commentary on AI and headcount. Investor briefing, 2026.

Wolf Street (2026). Hyperscaler capital expenditure analysis. 7 February 2026.

World Economic Forum (2025). *The Future of Jobs Report 2025*. WEF.

Zahra, S.A. and George, G. (2002). Absorptive capacity: A review, reconceptualization, and extension. *Academy of Management Review*, 27(2), 185-203.