

The United States as Outlier: Cross-Country Empirical Validation of the Access-Displacement Framework Across Ten Economies

Vito Henjoto

Antecedent Labs

ORCID: 0009-0003-0076-3665

March 2026

© 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/>

Companion paper: Henjoto, V. (2026). "Access Without Displacement: An Access-Displacement Framework for AI Economic Transformation." DOI: 10.5281/zenodo.19051765

Abstract

Henjoto (2026) proposed that technology-driven access and economic displacement are structurally independent, introducing the Access-Displacement Framework with seven testable contributions. That paper acknowledged a critical limitation: its empirical foundation was predominantly United States-centric. This companion paper addresses that limitation through systematic cross-country validation across ten economies spanning four income levels -- Japan, Germany, Switzerland, China, South Korea, Indonesia, India, Singapore, Taiwan, and Israel. Three empirical programmes are reported. First, the Access-Displacement validation confirms that the access gap ($G \sim A \sim$) and displacement gap ($G \sim D \sim$) operate independently in six economies with full confirmation and two additional cases with partial confirmation. China is the sole worker-level failure, under conditions that do not generalise, and the United States is identified as a global outlier rather than the representative case that displacement predictions assume. Five absorption mechanisms are taxonomised: demographic, institutional, flexible, segmented, and state-mediated. Financial sector "AI displacement" is systematically misattributed in all ten countries examined: every major banking "AI layoff" headline decomposes into mergers, regulatory consolidation, or demographic attrition when investigated. Second, the Measurement Obsolescence Hypothesis is validated internationally across nine countries, confirming that household labour force survey degradation is a global structural phenomenon, not a US-specific failure. Third, the Economic Forcing Function is quantified

empirically: frontier model training costs growing at 2.4 times per year while inference costs for equivalent capability decline 280-fold in two years, creating a measurable forcing function toward hybrid deployment architectures. Together, these three programmes provide cross-country empirical evidence for the Access-Displacement Framework's central claims.

Keywords: artificial intelligence, labour markets, displacement, cross-country evidence, measurement obsolescence, automation, economic forcing function

1. Introduction

Henjoto (2026)[[^]companion] proposed the Access-Displacement Framework, arguing that technology-driven capability access ($G \sim A \sim$) and economic displacement ($G \sim D \sim$) are structurally independent -- that the speed at which technology becomes accessible does not predict the speed at which it displaces workers. The paper advanced seven testable contributions, summarised here so that the present paper may be read independently:

1. AGI-C and AGI-R. The displacement tradition implicitly assumes that artificial general intelligence will *replace* human cognition across all economically relevant domains (AGI-R). The complementarian tradition implicitly assumes the opposite: that machine and human cognition will remain structurally complementary regardless of machine capability (AGI-C). Neither tradition names this assumption. Making the distinction explicit transforms an unstated framing commitment into a testable empirical question. Five converging lines of evidence -- from biological neuroscience, philosophy of mind, deployment data, experimental studies, and meta-analysis -- support AGI-C.

2. The Access-Displacement Framework. The *access gap* ($G \sim A \sim$) measures how quickly a technology becomes available; the *displacement gap* ($G \sim D \sim$) measures how quickly it restructures employment. Historically, access gaps have compressed from decades to years across successive technology generations, while displacement gaps have remained in the range of 15 to 40 years regardless of access speed. The two are structurally independent: forecasts that infer displacement timelines from access timelines commit a category error.

3. The Measurement Obsolescence Hypothesis (MOH). The establishment surveys (e.g. the US Current Employment Statistics) and household surveys (e.g. the Current Population Survey) on which displacement predictions depend are structurally degrading. US Bureau of Labor Statistics benchmark revisions escalated from historical norms to 3.4 times the historical average (March 2025) in the absence of any declared recession.

4. The Sampling Residualisation Hypothesis (SRH). Survey nonresponse is not random. Workers who have adapted to artificial intelligence (those in freelance, platform, and micro-enterprise arrangements) are systematically absent from household surveys, creating a dynamic bias that grows in proportion to the economic transformation being measured.

5. The Organisational Absorption Rate (OAR). Institutional capacity to integrate new technology operates as a binding rate constraint regardless of technical capability. Even JPMorgan Chase, with maximum capital, maximum CEO commitment, and explicit willingness to restructure, produced redeployment, not displacement. Klarna attempted to exceed the rate constraint and was forced to reverse.

6. The Economic Forcing Function (EFF). The death of Moore's Law and Dennard scaling creates an escalating cost ceiling on centralised computation. The industry has responded with miniaturisation and edge deployment, driving architecture toward hybrid human-AI systems rather than centralised replacement.

7. The Hybrid Swarm Architecture. The deployment topology that the Economic Forcing Function produces: centralised training, distributed edge inference, and a bidirectional integration layer connecting them. This is the architectural expression of AGI-C under cost constraints.

That paper carried a self-acknowledged limitation. Section 6.3 stated: "The empirical foundation of this paper is predominantly United States-centric. 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" (Henjoto, 2026, Section 6.3). This companion paper addresses that limitation directly.

The present paper reports three systematic empirical programmes conducted across multiple countries. The first programme validates the $G \sim A \sim G \sim D \sim$ independence claim across ten economies: Japan, Germany, Switzerland, China, South Korea, Indonesia, India, Singapore, Taiwan, and Israel. These ten countries were selected to span four income levels (high-income OECD, high-income non-OECD, upper-middle-income, lower-middle-income), three continents, and five distinct institutional configurations (coordinated market economies, liberal market economies, developmental states, pre-automation economies, and segmented dual economies). The second programme extends the Measurement Obsolescence Hypothesis from the United States to nine countries, testing whether household survey degradation is a structural global phenomenon or an artefact of specific American institutional failures. The third programme quantifies the Economic Forcing Function with empirical data on compute cost trajectories, hyperscaler capital expenditure, and the industry-wide pivot toward hybrid deployment architectures.

The central finding is that the United States -- source of the most influential automation-displacement research (Frey and Osborne, 2017; Acemoglu and Restrepo, 2020; Acemoglu, 2024; Autor, 2015) and the empirical basis for virtually all crisis-level displacement predictions -- is the global outlier, not the representative case. Guarascio, Piccirillo, and Reljic (2025), in a meta-analysis spanning multiple countries, found that the US is the only major economy where automation-employment studies consistently yield negative results. In six of the ten economies examined here, $G \sim A \sim$ and $G \sim D \sim$ operate independently with full confirmation, with partial confirmation in two additional cases and a single failure under

non-generalisable conditions. Across all testable economies, automation is associated with either net employment growth or zero net effect. Financial sector "AI displacement," the domain where artificial intelligence should produce the clearest labour substitution, is systematically misattributed in ten of ten countries. Household survey degradation is confirmed in all nine countries examined, operating through different mechanisms but producing the same structural outcome: declining reliability of the employment data on which displacement predictions depend. And the Economic Forcing Function is quantifiable: frontier training costs growing at 2.4 times per year (Cottier, Rahman, Fattorini, Maslej, Besiroglu, and Owen, 2024) while inference costs for equivalent capability decline 280-fold in two years (Stanford Institute for Human-Centered Artificial Intelligence, 2025; Epoch AI, 2025), creating measurable economic pressure toward hybrid architectures that structurally embed human participation.

These findings do not demonstrate that displacement is impossible or that labour market disruption is illusory. They demonstrate that the empirical evidence, when examined systematically across countries rather than extrapolated from a single national context, does not support the displacement narrative's scope or timeline. The binding constraints on automation's labour market impact are institutional, demographic, and economic, not technological. The task-based framework (Acemoglu and Restrepo, 2019) correctly identifies that automation simultaneously displaces workers from existing tasks and creates new tasks, but the cross-country evidence shows that institutional structure, not task composition, determines which effect dominates. Policy designed around US-centric displacement predictions may be systematically miscalibrated for the majority of the world's economies.

The remainder of this paper proceeds as follows. Section 2 presents the cross-country Access-Displacement evidence. Section 3 examines financial sector displacement misattribution as a dedicated empirical test. Section 4 extends the Measurement Obsolescence Hypothesis internationally. Section 5 quantifies the Economic Forcing Function. Section 6 develops the Organisational Absorption Rate into a cross-country taxonomy. Section 7 discusses implications, limitations, and directions for future work. Section 8 concludes.

2. Cross-Country Access-Displacement Evidence

2.1 Country Selection and Methodology

Ten countries were selected to maximise variation across institutional structures, income levels, demographic trajectories, and automation intensity. Each country was examined through a systematic programme of evidence gathering (typically three to six research passes per country), followed by independent verification of all major claims and synthesis into a standardised analytical framework. Cross-country comparative studies that synthesise econometric findings, institutional data, and labour market statistics across multiple national contexts face inherent verification challenges: source documents span

multiple languages, institutional terminology varies across countries, and secondary reporting of statistical findings frequently introduces errors. To address this, each country was examined through multiple independent research passes targeting different dimensions (e.g. demographic structure, institutional mechanisms, firm-level evidence, measurement infrastructure), with approximately 1,400 data points verified against primary sources across the programme.

The methodology was comparative-institutional rather than quantitative-parametric. The programme draws on existing econometric studies conducted within individual countries (Adachi, Kawaguchi, and Saito, 2024, for Japan; Dauth, Findeisen, Suedekum, and Woessner, 2021, for Germany; Yang, 2022, for Taiwan; Jeong and Jo, 2025, for South Korea; Cali and Presidente, 2025, for Indonesia; Giuntella, Lu, and Wang, 2025, for China) and synthesises their findings into a cross-country framework. This approach follows the comparative institutional tradition of Esping-Andersen (1990) and Hall and Soskice (2001) rather than the quantitative panel estimation tradition.

2.2 Master Classification

Table 1 presents the classification of all ten countries.

Table 1. Access-Displacement Classifications Across Ten Economies

Country	Classification	Primary Mechanism	$G\sim A\sim / G\sim D\sim$ Independent?	Robot Density (per 10K)	TFR	Key Distinguishing Feature
Japan	Demographic Absorption	Aging + firm-level internal labour markets (ILMs)	Yes	419	~1.20	Robots increase employment (Adachi et al., 2024; robot density: IFR, 2024)
Germany	Institutional Absorption / Vulnerable Periphery	Dual labour market + Mittelstand	Yes	429	~1.35	Institutional erosion: collective bargaining 80% to 49%
Switzerland	Flexible Absorption	Market flexibility + weak employment protection legislation (EPL)	Yes	296	~1.46	Displacement export via 403,000 cross-border workers
China	State-Mediated Bifurcation	State direction + development stage	Firm yes; worker no	470	~1.0	Only worker-level $G\sim A\sim / G\sim D\sim$ failure

South Korea	Demographic Absorption	Aging + chaebol ILMs	Yes	1,012	0.72	AI and robots diverge: AI +4.8%, robots -3.3%
Indonesia	Oligarchic Deferral	Pre-automation threshold	Yes	~5	~2.1	Trade competition, not technology, drives displacement
India	Stratified Deferral	Caste barrier + digital leapfrog	N/A	~7	~2.0	UPI/Jan Dhan: over 570 million bank accounts (early 2026)
Singapore	Institutional Absorption + 3 qualifiers	State-directed + supply-side	Yes	770	0.97	Supply-side absorption prevents displacement
Taiwan	Demographic Absorption + 2 qualifiers	Aging + semiconductor cluster	Yes	276	<0.8	TSMC 5.3x employment growth while most automated
Israel	Segmented Dual Economy + 2 qualifiers	Military-mediated $G\sim A\sim$ + demographic inversion	Yes (within tech)	Low (AI-first)	2.91 (2024)	Three-economy segmentation; AI-first automation

2.3 $G\sim A\sim/G\sim D\sim$ Independence: Cross-Country Evidence

The structural independence of the access gap and the displacement gap -- the central claim of the Access-Displacement Framework (Henjoto, 2026, Section 2.1) -- is confirmed across the programme. Six economies show full confirmation of $G\sim A\sim/G\sim D\sim$ independence (Japan, Germany, Switzerland, South Korea, Singapore, Taiwan). Two additional economies provide partial confirmation: Israel's independence holds within the tech sector (11.5 per cent of the workforce) but is untested in the broader economy; Indonesia shows positive employment effects at early-stage automation but operates below the adoption threshold where the framework's dynamics engage. India is excluded as not yet testable (below the automation threshold with no industrial-robot displacement to measure). China is the sole failure case, discussed separately below. Table 2 summarises the evidence.

Table 2. $G\sim A\sim/G\sim D\sim$ Independence Evidence by Country

Country	$G\sim A\sim$ Evidence	$G\sim D\sim$ Evidence	Independence	Source
---------	------------------------	------------------------	--------------	--------

			Mechanism	
Japan	Robots +2.2% employment	Manufacturing declining	Firm-level complementarity	Adachi, Kawaguchi, and Saito (2024)
Germany	Zero net effect	Atypical workers vulnerable	Sectoral: mfg $G \sim D \sim$ absorbed by services $G \sim A \sim$	Dauth, Findeisen, Suedekum, and Woessner (2021)
Switzerland	+800,000 net jobs	Banking -9,700 FTE/decade	Geographic: displacement exported	Deloitte Switzerland (2015)
South Korea	AI +4.8% employment	Robots -3.3% employment	Technology-type divergence	Jeong and Jo (2025)
Singapore	Employment grew continuously	DBS 4,000 contract cuts	Workforce-layer: permanent grows, flexible absorbs	MAS financial sector data
Taiwan	AI positive (firm-level)	Blue-collar -0.94%	Occupational: white-collar $G \sim A \sim$, blue-collar $G \sim D \sim$	Yang (2022)
Indonesia	Positive at low automation	Trade-driven layoffs	Development-stage: positive at early automation	Cali and Presidente (2025)
Israel	Tech sector 403,000 (11.5% of workforce), 2.8x wage premium	Admin -4.1%, HR - 8%	Occupational-tier: high-skill $G \sim A \sim$, routine $G \sim D \sim$	Israel Innovation Authority (2025a, 2025b)

China is the sole exception. Firm-level $G \sim A \sim / G \sim D \sim$ independence holds (BYD, CATL, and Freeman et al. (2025) documented +13.8 per cent employment growth alongside automation). But at the worker level, Giuntella, Lu, and Wang (2025), published in the *Economic Journal*, found that robot adoption reduced employment probability by 6.7 per cent. This occurs under specific conditions: a population exceeding one billion, unevenly distributed development, and rapid automation imposed through state direction. These conditions do not generalise. India and Indonesia are below the automation adoption threshold (approximately 5 to 10 robots per 10,000 manufacturing workers) and are therefore not testable for $G \sim A \sim / G \sim D \sim$ independence in the industrial-robot sense, though India's Unified Payments Interface demonstrates $G \sim A \sim$ in a different domain.

2.4 The US as Outlier

Guarascio, Piccirillo, and Reljic (2025), in a meta-analysis spanning multiple countries, found that the United States is the only major economy where automation-employment studies consistently find negative employment effects. Most countries show positive or neutral effects. Bessen (2019) documented the same pattern historically: automation has consistently boosted employment at the industry level, with displacement occurring primarily through cross-industry reallocation rather than net job loss. This finding reframes

the entire displacement debate. The automation-displacement literature is overwhelmingly US-centric: Frey and Osborne (2017) use US occupational data; Acemoglu and Restrepo (2020) use US commuting zones; Arntz, Gregory, and Zierahn (2016) recalibrated the Frey and Osborne estimates for OECD countries but retained the US task-based framework; Goldman Sachs (2023) extrapolates global estimates from US-derived task analysis. Even the foundational task-based framework (Autor, Levy, and Murnane, 2003) was developed from US occupational data. If the US is the outlier, then global displacement predictions built on US evidence inherit a systematic bias.

Why might the US be the outlier? I identify three features that distinguish the American labour market from the other economies examined here, though I acknowledge this is an interpretive claim that the comparative evidence suggests rather than proves. First, the US has the weakest institutional buffers among high-income countries: no works councils, minimal collective bargaining (approximately 6 per cent union density in the private sector), at-will employment in most states, and no equivalent to the German *Kurzarbeit*, Singapore SkillsFuture, or Japanese lifetime employment norms. Second, the US has the most financialised corporate governance: quarterly earnings pressure incentivises headcount reduction as a visible cost-cutting signal in ways that stakeholder-oriented systems (Germany, Japan) do not. Third, the US has the highest inequality among advanced economies, concentrating automation's benefits in narrow corridors without the redistributive mechanisms that other systems employ.

2.5 Five Absorption Mechanisms

The programme identifies five distinct mechanisms through which economies manage the automation-employment transition. Each is summarised here; Section 6 develops them into rate-constrained models with binding variables that determine absorption speed.

Demographic absorption (Japan, South Korea, Taiwan): population decline and rapid aging create structural labour scarcity that absorbs automation's displacement potential. Acemoglu and Restrepo (2022) demonstrated that aging explains approximately 40 per cent of cross-country variation in robot adoption.

Institutional absorption (Germany, Singapore): works councils, collective bargaining, tripartite systems, and state-directed retraining convert involuntary layoffs into managed attrition, retraining, and redeployment. The limitation is institutional erosion: Germany's collective bargaining coverage has declined from approximately 80 per cent (early 1990s) to approximately 49 per cent (2023).

Flexible absorption (Switzerland): low employment protection (OECD EPL index 2.10) combined with high market flexibility allows rapid sectoral reallocation, partly dependent on displacement export through 403,000 cross-border workers.

Segmented dual economy (Israel): the automation economy and the demographic growth economy operate in different population segments. Technology neither displaces nor benefits the fastest-growing populations (Haredi: 13.9 per cent of population, projected 32

per cent by 2065; Israel Democracy Institute, 2024, 2025). Israel ranks first globally on the Anthropoc Economic Index (7.0; Anthropic, 2025) and is the first economy in the programme where cognitive automation rather than industrial robots is the primary vector. The Israel Defence Forces (IDF) pipeline -- particularly Unit 8200 (signals intelligence), producing approximately 50 per cent of founders whose companies were acquired for over \$100 million (Swed and Butler, 2015) -- functions as a structured access-creation mechanism that channels approximately 1,000 to 1,500 graduates per year into the civilian tech sector.

State-mediated bifurcation (China): state direction simultaneously promotes automation and manages displacement through administrative means, producing the only worker-level $G\sim A\sim/G\sim D\sim$ failure in the programme under population-scale conditions that do not generalise.

Indonesia and India are below the automation adoption threshold. Their displacement dynamics are driven by trade competition (Indonesia: 77,965 total layoffs in 2024, with Chinese import competition the dominant driver in manufacturing; Winters, 2011), oligarchic rent extraction, and caste-mediated stratification, not by technology. India's UPI/Jan Dhan revolution (172 billion transactions and over 570 million bank accounts for the previously unbanked) demonstrates technology operating as an access-creation mechanism rather than a displacement mechanism.

2.6 Country Evidence Summaries

Japan. Japan's labour market absorbs automation through demographic replacement rather than institutional negotiation. The population declined by approximately 908,000 in 2024 alone, unemployment stands at approximately 2.5 per cent (a historic low), and robot density is the third-highest in Asia at 419 per 10,000 manufacturing workers (IFR, 2024). The mandatory retirement system (age 60 or 65, with re-hiring at reduced wages) creates a predictable vacancy flow that automation fills without displacing incumbent workers. The megabank pattern illustrates the mechanism: Mizuho announced 19,000 and MUFG 9,500 "AI-driven" position reductions, but both specified attrition of bubble-era retirees, not layoffs. MUFG's subsequent trajectory, examined in detail in Section 3.3, illustrates the point: the bank grew 24 per cent after announcing AI cuts. What media frames as "AI displacement" is routine generational turnover operating through demographic absorption channels. The demographic absorption mechanism is quantified by Adachi, Kawaguchi, and Saito (2024), who found that a one-standard-deviation increase in robot adoption raised establishment-level employment by 2.2 per cent, and corroborated by Ni and Obashi (2021) at the firm level, where robot adoption increases both job creation and job destruction with net effects varying by firm size and sector. Kikuchi, Fujiwara, and Shirota (2024) confirmed that while automation reduces routine occupations in Japan, it does not reduce overall employment -- workers transition to non-routine tasks rather than unemployment. In labour-scarce sectors, the effect inverts: Lee, Iizuka, and Eggleston (2025) found that robot adoption in Japanese nursing homes is associated with increased employment and improved staff retention, with automation operating as a recruitment tool.

Henn-na Hotel (which fired half its robots in 2019) and SoftBank Pepper (production halted in 2021) demonstrate that service-sector automation of complex tasks remains unreliable.

Germany. Germany absorbs automation through institutional machinery (works councils, collective bargaining, Kurzarbeit, and the Mittelstand apprenticeship system) that converts involuntary displacement into managed attrition, retraining, and redeployment. The net employment effect of robot adoption is zero: manufacturing losses are fully offset by service sector gains, as Dauth, Findeisen, Suedekum, and Woessner (2021) demonstrated in the *Journal of the European Economic Association*. This zero-net-effect finding is consistent with Koch, Manuylov, and Smolka (2021), who found positive firm-level employment effects from robot adoption using Spanish data, and with Oesch and Piccitto (2019), who showed that German occupational structure underwent upgrading, not polarisation, between 1992 and 2015. The historical depth of Germany's adaptive capacity is documented by Spitz-Oener (2006): over 99 per cent of West German workers adjusted their task composition between 1979 and 1999, with fewer than 1 per cent losing employment to automation. But this institutional architecture is eroding. Collective bargaining coverage has declined from approximately 80 per cent in the early 1990s to approximately 49 per cent in 2023, with two-thirds of private sector workers now lacking any form of co-determination (Kohaut and Schnabel, 2024). Works councils cover only approximately 42 per cent of West German and 37 per cent of East German employees. The Hamburg Labour Court (2024) created a co-determination gap for individual AI tool use through private accounts. The erosion is consequential because AI affects the workplace differently from industrial robots: Gathmann, Grimm, and Winkler (2024) found that AI reduces non-routine abstract tasks while increasing demand for high-level routine tasks (the opposite of the pattern robots produce), while Dengler and Matthes (2018) documented that the share of German occupations with high substitution potential rose from 15 per cent in 2013 to 25 per cent in 2016. Germany thus presents a dual finding: institutional absorption works, but the institutions doing the absorbing are themselves in structural decline. Allianz Partners announced 1,500 to 1,800 genuine AI-driven job cuts in call centres, the only clean AI displacement case in the entire ten-country programme. Yet Destatis data shows aggregate financial sector employment grew by 12,000 in 2024 and 16,000 in 2025: even genuine firm-level displacement is absorbed at the sector level.

Switzerland. Switzerland absorbs automation through market flexibility and geographic displacement export rather than institutional negotiation. The country achieves the highest employment rate in Europe (80.7 per cent) with the weakest employment protection legislation (OECD EPL index 2.10), and approximately 800,000 net new jobs were created between 1990 and 2013 despite continuous automation (Deloitte Switzerland, 2015). The mechanism combines two features. First, low employment protection allows rapid sectoral reallocation: workers move from declining to growing sectors within months rather than the years that institutional negotiation requires. Second, 403,000 cross-border workers from France, Germany, and Italy absorb Swiss displacement pressure, commuting daily to work in Swiss firms while living in lower-cost jurisdictions -- a geographic buffer that exports displacement to neighbouring economies. The interaction between immigration and automation in this flexible context is complementary, not

substitutive: Beerli, Ruffner, Siegenthaler, and Peri (2021), in a study published in the *American Economic Review*, found that the abolition of immigration restrictions between Switzerland and the EU caused native wages to increase rather than decline. Banking employment declined by approximately 9,700 full-time equivalents over the decade (2014 to 2024), but this is structural (ATM, online banking, then fintech) and predates AI. The UBS-Credit Suisse merger (2023) produced 36,000 cuts globally through merger consolidation, not automation. The broader financial services sector grew by 5.8 per cent over the same period as narrow banking declined -- displacement within banking was absorbed by adjacent financial services growth.

China. China is the sole economy in the programme where worker-level $G \sim A \sim G \sim D \sim$ independence fails, and the conditions that produce this failure do not generalise. The divergence between firm-level and worker-level outcomes is the defining feature. At the firm level, $G \sim A \sim G \sim D \sim$ independence holds: BYD, CATL, and other manufacturers grew employment while automating, and Freeman et al. (2025), studying 381,811 Chinese manufacturing firms, found that minimum wage increases drove robot adoption, which in turn increased employment by 13.8 per cent and wages by 13.0 per cent. But at the worker level, Giuntella, Lu, and Wang (2025), published in the *Economic Journal*, found that robot adoption reduced employment probability by 6.7 per cent. The firm creates jobs; the displaced worker cannot access them. This divergence operates through conditions specific to China: a population of 1.4 billion with extreme development unevenness (Shenzhen's GDP per capita exceeds \$30,000; rural Gansu is below \$5,000), rapid state-directed automation (Made in China 2025), and administrative displacement management (hukou restrictions, SOE employment guarantees, rural bank consolidation with 350 banking licences cancelled in 2025 alone, up from 198 in 2024). The human capital constraint is binding: Rozelle and Hell (2020) documented that only approximately 30 per cent of China's population has completed high school, creating a "displacement trap" where low-skilled workers cannot transition to the complementary roles that automation creates in high-income economies. The subjective experience of automation reflects this structural position: Wu and Sun (2025), in *The China Quarterly*, found that internal migrants with exit options (rural land holdings, geographic mobility) viewed automation with less anxiety than local workers -- "little to lose" -- because their alternative to factory work was not unemployment but return to a rural holding. The population-scale hypothesis, that $G \sim A \sim G \sim D \sim$ independence fails when an economy is simultaneously very large, unevenly developed, and automating rapidly under state direction, has $N=1$. India is the theoretical second candidate but is decades from China's automation level. I am not confident that the conditions specified here are exhaustive; the failure may involve mechanisms not yet identified, and the $N=1$ problem means this hypothesis cannot be tested further until another economy of comparable scale reaches comparable automation intensity.

South Korea. South Korea combines the world's highest robot density (1,012 per 10,000 manufacturing workers) with near-record employment and the world's lowest fertility rate (0.72 in 2023, until Taiwan's fell below 0.8 in 2025), creating a demographic absorption dynamic comparable to Japan's. The country's distinguishing contribution to the

programme is a technology-type divergence: AI and industrial robots have opposite employment effects within the same economy. Jeong and Jo (2025), in the *Weizenbaum Journal of the Digital Society*, found that AI adoption increases employment by 4.8 per cent while robot adoption decreases it by 3.3 per cent. This divergence, if it replicates across countries, would fundamentally alter the displacement debate, which has predominantly studied industrial robots. The chaebol structure (Samsung, Hyundai, SK, LG) provides internal labour markets for cross-subsidiary redeployment that cushion the robot-driven displacement. Exposure estimates have not translated into observed displacement: the Bank of Korea estimates 3.14 million jobs are highly exposed to AI replacement -- approximately 12 per cent of total employment at the time of the estimate (Han and Oh, 2025) -- but Park and Shin (2025), using propensity score matching, found that both AI and robot adoption are associated with increased employment in Korean firms. The IMF (2025) found that approximately 50 per cent of South Korean jobs are exposed to AI, split roughly evenly between high-complementarity roles (24 per cent) and low-complementarity roles (27 per cent) -- confirming that exposure alone does not predict whether AI will augment or substitute.

Indonesia. Indonesia operates below the automation adoption threshold (robot density approximately 5 per 10,000) where $G \sim A \sim G \sim D \sim$ dynamics engage, establishing the development-stage scope condition for the Access-Displacement Framework. Employment is growing (unemployment fell from 7.07 per cent in 2020 to 4.76 per cent in 2025), and the displacement that exists is trade-driven, not technology-driven: of the 77,965 total layoffs recorded in 2024 (Indonesian Manpower Ministry), Chinese import competition was the dominant driver in manufacturing. The 59.4 per cent informal sector represents a measurement black hole that makes any displacement assessment incomplete. Banking sector losses of 48,309 employees (2014 to 2023) were originally characterised as technology-driven, but decomposition revealed that Bank Danamon's MUFG merger accounted for 72 per cent of the total: merger consolidation, not automation. Chang and Huynh (2016) estimated that 56 per cent of ASEAN employment faces high automation risk, but this exposure has not translated into observed displacement at Indonesia's current automation levels. The positive employment effects at early-stage automation documented by Cali and Presidente (2025), with diminishing returns at higher levels, are consistent with the pre-automation scope condition: the A-D framework applies above a minimum automation density threshold, and Indonesia has not yet crossed it.

India. India (robot density approximately 7 per 10,000; IFR, 2024) is below the automation threshold for industrial robotics but demonstrates technology operating as an access-creation mechanism rather than a displacement mechanism. The Unified Payments Interface processed 172 billion transactions in 2024 (growing 46 per cent year-on-year), moving approximately \$2.9 trillion. The Pradhan Mantri Jan Dhan Yojana created over 570 million bank accounts for the previously unbanked as of early 2026 (approximately 56 per cent held by women). Technology here is not substituting existing workers; it is creating economic participation for populations that were structurally excluded. The fintech sector has grown rapidly, with over 9,000 startups and 26 unicorns at an industry valuation of approximately \$90 billion (2024), and traditional banking added 123,000 new positions in

financial year 2023 (a decade high), driven by expansion into tier-2 and tier-3 cities. The Paytm layoffs (4,600 workers in financial year 2025) were driven by Reserve Bank of India regulatory enforcement, not AI. India also reveals three scope conditions for the broader framework: caste barriers that stratify technology access independently of skill or education; the brain drain paradox (the IT sector employs millions but is an exception within the broader economy, not representative of it); and the digital leapfrog phenomenon (UPI bypasses card infrastructure entirely, creating a methodological boundary for robot-employment literature that assumes manufacturing as the primary automation channel).

Singapore. Singapore operates a supply-side absorption model that differs fundamentally from the demand-side models observed in Germany and Switzerland: rather than managing displacement after it occurs, Singapore prevents displacement by constraining labour supply. Immigration policy, compulsory National Service (24 months, delaying male labour market entry), and demographic pressure (total fertility rate 0.97 in 2023-2024, a historic low) limit the supply of workers competing with automated systems, while the state-directed SkillsFuture system and tripartite coordination reskill the existing workforce. The result is simultaneous AI deployment and employment growth: the financial sector added an average of 4,400 net jobs per year from 2021 to 2024, exceeding the Industry Transformation Map target. DBS Bank illustrates the mechanism at the firm level. Its then-CEO Piyush Gupta stated in February 2025 that "for the first time, I'm struggling to create jobs." Yet the bank deployed 370 AI use cases powered by over 1,500 models and reported S\$1 billion in AI-derived economic value. The 4,000 job cuts targeted contract and temporary staff; permanent headcount grew 2.5 per cent to 41,354. The displacement was absorbed through the workforce's flexible layer while the permanent layer expanded. Singapore ranks second globally on the Anthropic AI Usage Index (4.57) after Israel (7.0), confirming that high AI adoption is compatible with growing employment when institutional absorption mechanisms are in place.

Taiwan. Taiwan demonstrates demographic absorption at its most extreme, combined with the most pronounced dual economy in the programme outside Israel. The total fertility rate fell below 0.8 in 2025 -- the world's lowest, surpassing South Korea -- and Taiwan became a super-aged society (20 per cent aged 65 and over) in December 2025, completing the transition from aged to super-aged in seven years (compared with 11 for Japan, 19 for Italy, and 36 for Germany). The long-term macroeconomic consequences are modelled by Cheng and Loichinger (2017) and Goh, Wong, McNown, and Chen (2023). The semiconductor sector illustrates how automation and employment growth coexist under demographic pressure: TSMC grew from 15,880 employees in 2000 to 83,825 in 2024 -- a 5.3-fold increase during the most automation-intensive period in semiconductor manufacturing history, with annual capital expenditure averaging approximately TWD 1 trillion. At the firm level, Yang (2022), in *Research Policy*, found that AI technology is positively associated with both productivity and employment in Taiwan's electronics industry. But this growth is concentrated in a narrow corridor. The Silicon Valley-Hsinchu technical community connection documented by Saxenian and Hsu (2001) has produced a dual economy in which the semiconductor hub's median annual salary is 65 per cent above the national median, electronics manufacturing accounts for more than 15 per cent

of GDP from only 6.5 per cent of the workforce (Capital Economics, via CNN, November 2025), and real wages for the broader economy stagnated from approximately 2000 to 2020. Genuine wage growth returned only in 2024 to 2025 as demographic tightening forced employer competition for workers -- the labour scarcity mechanism finally reaching beyond the semiconductor corridor.

Israel. Israel inverts the demographic logic of every other economy in the programme: it faces labour surplus, not scarcity, with a fertility rate of 2.91 in 2024 (the highest in the OECD, nearly double the average) and a growing population of 10.2 million. But the growth is concentrated in the populations with lowest technology participation, producing a three-economy segmentation. The Haredi (ultra-Orthodox) community, comprising 13.9 per cent of the population (projected 32 per cent by 2065; Israel Democracy Institute, 2024, 2025), has male labour force participation of 54 per cent and tech participation of 3.5 per cent. The Arab-Israeli population at 21.1 per cent has tech participation of 1.5 per cent. The secular tech economy (403,000 workers, 11.5 per cent of the workforce) generates 17.3 per cent of GDP and 53 to 57 per cent of exports at a wage premium of 2.8 times the national average (Israel Innovation Authority, 2025a, 2025b), with R&D spending at 6.35 per cent of GDP -- the highest in the world. The OECD projects a 12-percentage-point GDP per capita gap between full integration and frozen participation rates by 2065 (OECD, 2025). Technology neither displaces nor benefits the fastest-growing populations; it operates in a segment they do not participate in. Israel is also the first economy in the programme where AI rather than industrial robots is the primary automation vector, and the Taub Center (Debowy, Winter, Epstein, Weiss, and Behar-Netanel, 2025) found that 9 per cent of AI-adopting businesses report reducing their workforce -- an early signal that AI displacement may operate through different mechanisms than robot displacement. The tech sector experienced its first employment decline in a decade in 2024 (-5,000 net) and R&D employment fell 6.5 per cent in the first half of 2025, though both figures are confounded by the October 7 war (360,000 reservists authorised, approximately 287,000 mobilised; net citizen emigration of 125,200 between 2022 and August 2024).

2.7 Distributional Displacement: A Novel Extension

The country-level evidence establishes $G \sim A \sim / G \sim D \sim$ independence at the aggregate level. But two countries -- Taiwan and Israel -- reveal a pattern that the Access-Displacement Framework as formulated in Henjoto (2026) did not anticipate, and which this paper proposes as a novel extension.

In both economies, aggregate employment grows while automation concentrates its benefits in narrow corridors. Taiwan's Hsinchu semiconductor corridor earns 65 per cent above the national median; electronics generates 15 per cent of GDP from 6.5 per cent of the workforce. Israel's tech sector commands a 2.8-times wage premium from 11.5 per cent of workers. South Korea's chaebols provide internal redeployment for core employees while platform workers absorb precarity. In each case, workers are not losing jobs -- they are excluded from productivity gains. I term this **distributional displacement**: technology

concentrating productivity gains in narrow geographic, demographic, or institutional corridors without producing aggregate job loss.

Distributional displacement is distinct from three established concepts in the inequality literature with which it shares surface-level similarities. It is not **labour market polarisation** (Autor, 2015; Goos and Manning, 2007), which describes the disappearance of middle-skill occupations as workers are pushed toward high-skill and low-skill poles. Polarisation involves occupational restructuring -- jobs are lost in the middle of the distribution. Distributional displacement involves no occupational restructuring. The occupational structure remains stable; the income geography changes. It is not **skill-biased technological change** (Goldin and Katz, 2008), which assumes a meritocratic continuum in which more education and higher skill yield greater returns from technology. Distributional displacement operates across structural binary boundaries that are not skill-based: the boundary in Taiwan is geographic (Hsinchu versus the rest of the island), in Israel it is demographic-institutional (the IDF pipeline versus excluded populations), and in South Korea it is organisational (chaebol core versus SME/platform periphery). A skilled Arab-Israeli computer science graduate faces a 1.5 per cent tech participation rate not because of skill deficiency but because of structural exclusion from the military pipeline that produces the majority of acquired-company founders (Swed and Butler, 2015). Education and upskilling -- the standard policy response to SBTC -- are necessary but insufficient when the boundary is structural rather than meritocratic. And it is not **capital-income divergence** (Piketty, 2014), which describes returns to capital exceeding returns to labour. Distributional displacement is a within-labour phenomenon: tech-participating labour versus non-tech-participating labour, separated by structural boundaries, not by asset ownership.

Brynjolfsson (2022) identified the "Turing Trap" at the firm and industry level: AI systems designed to replicate human performance concentrate returns among system owners rather than augmenting broad-based productivity. The Taiwan and Israel evidence suggests this trap is already operating at the national level, producing economy-wide segmentation through the geographic and demographic concentration of automation's benefits.

This finding extends the Access-Displacement Framework beyond its original formulation. Henjoto (2026) defined $G\sim D\sim$ as the distance between current workforce levels and the level that would obtain if all technically feasible displacement were realised -- an employment-count measure. The Taiwan and Israel evidence suggests that $G\sim D\sim$ requires a distributional component alongside its employment component: automation can produce a dual economy in which the displacement gap remains closed (no net job loss) while the welfare gap widens (productivity gains accrue to a narrow segment). This extension may explain why the displacement narrative persists despite favourable aggregate employment data: people experience distributional exclusion, which feels like displacement even when employment statistics say otherwise.

2.8 Immigration-Automation Substitution

A second cross-cutting pattern emerged across all high-income economies: migrant labour defers automation adoption. Every country with significant automation uses foreign workers to fill positions that might otherwise drive technology investment, creating a buffer between capability access and displacement realisation.

Israel's post-October 7, 2023, response illustrates the mechanism directly. Following the revocation of approximately 165,000 Palestinian work permits (reduced to approximately 44,200 by June 2025), Israel (home to 170 robotics companies, ranking first globally in AI adoption per capita) replaced the labour gap not with automation but with Indian construction workers (approximately 16,000 in 2024 under a bilateral agreement). Taiwan employs 866,000 blue-collar migrant workers (7.4 per cent of the workforce), including approximately 480,000 in manufacturing (17.6 per cent of manufacturing employment) and approximately 75 per cent of the formal care workforce, a long-standing pattern of labour importation substituting for automation investment (Tsay and Lin, 2001). Singapore's foreign worker system (approximately 1.4 million, or 37 per cent of the workforce) functions as a demographic buffer. Switzerland exports displacement through 403,000 cross-border workers.

This pattern suggests that the pathway from capability access to labour displacement runs through immigration policy as an intermediate step. In the cases examined, automation became the preferred response only when migrant labour was unavailable, unaffordable, or politically constrained. The displacement timeline is therefore governed partly by immigration dynamics, a determinant absent from most automation-employment models. Whether this represents a genuine causal mechanism or a correlation driven by third factors (labour market tightness, demographic structure) that independently affect both immigration policy and automation adoption is not resolved by the evidence presented here.

3. Financial Sector Displacement Misattribution

3.1 The Test

If artificial intelligence were going to displace workers at scale, it should happen first in financial services. Banking involves routine cognitive tasks (transaction processing, compliance checking, customer queries), operates on massive structured datasets, and offers a clear business case for automation (labour is typically 40 to 60 per cent of operating costs). Every major AI-displacement narrative includes banking as a leading indicator. The financial sector therefore provides a natural test: if systematic AI-driven displacement is not occurring in banking, the aggregate displacement thesis faces a significant empirical challenge.

3.2 Results: Ten of Ten Countries

Financial sector "AI displacement" was examined across all ten countries. In every case, major "AI layoff" headlines decomposed into non-AI drivers when investigated. Table 3 summarises the findings.

Table 3. Financial Sector Displacement Decomposition

Country	Headline	Actual Primary Driver	Genuine AI?
Japan	Megabanks "cut 33,000" (2017)	Demographic attrition (bubble-era retirements)	No. MUFG grew 24% after announcing AI cuts.
Germany	Deutsche Bank 18,000 (2019)	Failed Commerzbank merger; investment banking retreat	No (2019). Sector grew +16,000 in 2025.
Germany	Allianz Partners 1,500-1,800	AI call centre automation	Yes -- sole genuine case in programme
Switzerland	UBS-Credit Suisse 36,000+	Merger consolidation	No
Singapore	DBS 4,000	Contract/temporary staff only; permanent headcount +2.5%	Partial -- flexible layer only
South Korea	~64,000 decline (2013-2025)	Branch closures + voluntary retirement packages	No
China	Big Four declining 1-2%/year	State-directed rural bank mergers (350 licences cancelled in 2025, up from 198 in 2024)	No
India	Paytm 4,600	Reserve Bank of India regulatory enforcement	No. Sector added 123,000 in FY2023 (decade high).
Indonesia	-48,309 (2014-2023)	Bank Danamon MUFG merger = 72% of total	No
Taiwan	No significant headlines	Demographic labour shortage prevents displacement	N/A
Israel	Bank Hapoalim 770	Managed voluntary attrition	No

3.3 The MUFG Case

MUFG (Mitsubishi UFJ Financial Group), Japan's largest bank, illustrates the misattribution pattern in detail. In 2017, MUFG announced plans to automate 9,500 clerical positions through artificial intelligence. Media coverage framed this as AI replacing bank workers. MUFG's actual headcount trajectory: 121,800 employees (2023) growing to 150,800 (2025),

a 24 per cent increase after announcing "AI job cuts." The bank now employs substantially more people than it did before the AI announcement. The 9,500 "AI cuts" were absorbed through demographic attrition (bubble-era hires reaching mandatory retirement age) while the bank simultaneously expanded in other functions.

3.4 The Allianz Exception and the Klarna Reversal

Allianz Partners (Germany) represents the sole genuinely AI-driven displacement case across ten countries: 1,500 to 1,800 call centre jobs explicitly attributed to AI automation of simple insurance queries, announced November 2025. This is real. It is also a call centre subsidiary handling routine queries, the narrowest and most automatable category of financial services work.

The Klarna case provides an instructive complement. Klarna (Sweden), a financial technology company, attempted pure AI displacement in customer service, reducing headcount from approximately 3,000 to approximately 2,000 agents. In mid-2025, the company reversed course: the CEO admitted quality had collapsed and Klarna began rehiring human agents in a hybrid model. This is the only major case in the programme where AI-driven displacement was attempted at scale, failed, and was explicitly reversed.

3.5 Displacement Attribution Decomposition

Across approximately fifteen major financial sector "AI displacement" events examined in the programme, the actual drivers decompose as follows:

- Merger and acquisition consolidation: approximately 40 to 50 per cent of announced cuts
- Demographic attrition (not replacing retirees): approximately 20 to 30 per cent
- Branch rationalisation and digital channel migration: approximately 10 to 15 per cent
- Regulatory and political restructuring: approximately 5 to 10 per cent
- Genuine AI-driven displacement: approximately 5 to 10 per cent

The headline narrative -- "AI is killing banking jobs" -- systematically misattributes the cause. The displacement is real but the attribution is wrong. This inflates perceived $G\sim D$ across all countries and sectors: if the sector most favourable to AI displacement shows 90 per cent misattribution, other sectors with weaker automation business cases likely show higher misattribution still.

4. International Measurement Obsolescence

4.1 Extending MOH Beyond the United States

Henjoto (2026, Section 3.1) documented accelerating degradation in US Bureau of Labor Statistics benchmark data, with revisions escalating from historical norms to 3.4 times the historical average. The Measurement Obsolescence Hypothesis proposed that this represented structural measurement failure, not transient estimation error. The hypothesis generated four testable predictions: (a) declining response rates, (b) growing survey-

administrative divergence, (c) systematic under-coverage of non-traditional populations, and (d) institutional acknowledgment of measurement limitations. This section tests those predictions across nine countries.

4.2 Nine-Country MOH Classifications

Table 4 presents the MOH classification for each country examined.

Table 4. Measurement Obsolescence Classifications

Country	Classification	Core Evidence
United Kingdom	Collapsed	LFS suspended (2023); de-accredited; response rate 79% to 21%
United States	Advanced Degradation	20pp response decline; 862,000 benchmark revision (3.4x historical)
Canada	Structural Decline	22pp response decline; 1 million temporary resident undercount
Italy	Accelerating Decay	Non-response +60% in four years; 3.1 million irregular FTEs
Germany	Structural Disruption	2020 overhaul created permanent break; non-response 3.4% to 46.8% to 9.1%
France	Slow Erosion	Compulsory but 2.2 million worker gap; 24.1% non-response
Spain	Structural Concealment	Low headline non-response masks ~49% proxy rate; 831,000 hidden unemployed
Japan	Institutional Opacity	Compulsory but zero response rates published; zero statistical corrections
Australia	Resilient (biased frame)	93% response rate but biased sampling frame; no ethnic employment data

The classifications are not arbitrary labels. They describe a degradation spectrum from total collapse (United Kingdom) to maintained but biased resilience (Australia).

The **United Kingdom** represents the endpoint of voluntary survey degradation. The Labour Force Survey response rate fell from approximately 79 per cent (1993) to a nadir of 12.7 per cent (2023), at which point the UK Statistics Authority de-accredited the survey's National Statistics designation. ONS suspended the LFS entirely and commissioned an independent review (Devereux, 2025), which found that most problems with core economic statistics were "the consequence of ONS's own performance" resulting from "inadequacies in the way ONS has made decisions, planned and budgeted, and managed

risks." Corlett and Slaughter (2024) documented the consequences for UK economic policymaking of operating without reliable labour market data. The National Statistician called publicly for compulsory participation -- an implicit admission that the voluntary model has structurally failed. A replacement survey is in development but reliable monthly UK employment data does not currently exist.

The **United States** shows advanced degradation documented in detail in Henjoto (2026, Section 3). CES benchmark revisions escalated from -187,000 (March 2023, within historical norms) to -598,000 (March 2024, 2.3 times the historical average) to -862,000 (March 2025, 3.4 times the average) -- the largest final benchmark revision since the Great Recession, in the absence of any declared recession. The CPS response rate reached a series low of 64 per cent in November 2025. Bracha and Burke (2023), at the Federal Reserve Bank of Boston, estimated that official employment statistics understate actual employment by 0.25 to 5.1 percentage points due to uncounted informal and gig work, while Abraham, Haltiwanger, Sandusky, and Spletzer (2021) documented systematic misclassification between IRS and CPS data on self-employment. A 43-day government shutdown (October-November 2025) rendered October household survey data unrecoverable.

Canada experienced a 22-percentage-point response rate decline (approximately 95 per cent in 1995 to approximately 73 per cent in 2025). StatCan's 2019 internal audit found that "management does not have a full understanding of the impact of non-response on statistical biases." The documented one-million temporary resident undercount (IRCC administrative data versus LFS estimates, December 2022) demonstrates that Canada's measurement gap is concentrated in precisely the population most relevant to immigration-automation dynamics.

Italy shows the fastest rate of deterioration. Non-response increased 60 per cent in four years (14.4 per cent to 23.1 per cent, 2018 to 2022) despite compulsory participation. Non-contacts account for 76.6 per cent of all non-response -- ISTAT literally cannot reach the respondents, regardless of legal obligation. Approximately 3.1 million irregular full-time equivalent workers (ISTAT's own estimate) are invisible to the survey. The employment gap between LFS and administrative sources stands at approximately -1.9 million (-7.7 per cent), with the gap widening rather than narrowing.

Germany engineered a methodological break. The 2020 Microcensus overhaul -- nominally a "modernisation" -- produced non-response rates that surged from 3.4 per cent (pre-2018) to 46.8 per cent (2020), settling to a new normal of 7 to 12 per cent (approximately three times the pre-2018 baseline). The German-language-only questionnaire excludes effective participation from 20.9 per cent of the population that is foreign-born. Schneider (2022) estimated Germany's shadow economy at EUR 481 billion (2024), representing an entire stratum of economic activity invisible to household surveys. Computer-Assisted Personal Interviewing (CAPI) collapsed to 5.1 per cent of interviews. Two consecutive census enumerations produced overcounts of approximately 1.4 to 1.5 million relative to register-based population data.

France maintains compulsory participation with nominal response rates near 76 per cent, but a 2.2-million worker gap between the Enquete Emploi and administrative employment sources reveals that the survey is measuring a shrinking share of the actual labour market. Non-response reached 24.1 per cent, concentrated in non-standard housing and immigrant populations.

Spain presents structural concealment: low headline non-response (15.7 per cent) masks a proxy response rate of approximately 49 per cent -- nearly half of all responses are provided by someone other than the person whose employment status is being recorded. Additionally, 831,000 fijo-discontinuo (fixed-discontinuous contract) workers are statistically invisible: reclassified from temporary to permanent contracts in 2022, they are counted as employed during inactive periods, artificially inflating the employment count.

Japan is unique in refusing to publish response rates at all -- the only G7 country where measurement quality cannot be externally assessed. The Statistics Bureau applies zero statistical corrections to the Labour Force Survey (triple-confirmed against ILO standards). Approximately 1.18 million workers are classified "unknown industry" and 980,000 "unknown occupation" -- all left unimputed. Japan has the institutional capability to impute (demonstrated in the Economic Census) but deliberately does not apply it to the LFS. The ILO's quality assessment (SSM3) for Japan contains seven or more errors, and the IMF's Data Quality Assessment Framework was last applied in 2006 -- predating both the Funabashi statistical fraud scandal (2007) and the Maetsu scandal (2018).

Australia serves as a partial control case: the only country maintaining response rates above 90 per cent, owing to compulsory participation with enforcement. However, the Australian Bureau of Statistics collects no racial or ethnic employment data (a structural limitation unique among OECD countries), cannot track employment outcomes by individual country of birth at monthly frequency (only at five-yearly Census intervals), and has experienced a 20 per cent decline in real resources over the decade to 2018. The 93 per cent response rate may reflect legal compulsion overriding the structural trend rather than the absence of the structural trend.

4.3 All Four Predictions Confirmed

Prediction (a): Declining response rates. Confirmed in all nine countries. The decline is structural and multi-decade, not cyclical or pandemic-related: de Leeuw, Hox, and Luiten (2018) documented the pattern across 36 years of Labour Force Survey data internationally. The consequences extend beyond data quality to active distortion of the macroeconomic indicators on which policy depends: Bernhardt, Munro, and Wolcott (2024) estimated that nonresponse alone accounts for at least 10 per cent of the reported decline in US labour force participation since 2000. Leduc, Oliveira, and Paulson (2025), writing for the Federal Reserve Bank of San Francisco, offered a partial counterpoint, finding that recent data revisions in payroll employment and CPI inflation have remained in line with pre-pandemic averages -- though they acknowledged that the structural decline in response rates remains a concern for future data quality. The United Kingdom collapsed from

approximately 79 per cent to approximately 21 per cent (with a nadir of 12.7 per cent in 2023). The United States Current Population Survey reached a series low of 64 per cent in November 2025. Canada declined from approximately 95 per cent to approximately 73 per cent. Even compulsory-participation countries showed decline: Italy's non-response increased 60 per cent in four years, and Germany's 2020 overhaul produced a 46.8 per cent non-response rate.

Prediction (b): Growing survey-administrative divergence. Confirmed in every country where administrative data exists for comparison. The United States: 862,000 benchmark revision. The United Kingdom: approximately one million worker divergence. Canada: 1,082,585 temporary resident undercount. France: 2.2 million worker gap. Italy: 3.1 million irregular full-time equivalents. Germany: divergence widening, with two consecutive census overcounts of approximately 1.4 to 1.5 million.

Prediction (c): Systematic under-coverage of non-traditional populations. Confirmed across all nine countries with consistent patterns. The same populations appear in under-covered lists across countries with independently designed surveys: young adults (18 to 34), foreign-born and ethnic minorities, gig and platform workers, and those in non-standard housing.

Prediction (d): Institutional acknowledgment. Confirmed across the full spectrum. The United Kingdom: survey suspended and de-accredited; the National Statistician called publicly for compulsory participation. The United States: Congressional Research Service documented the benchmark revision process amid growing concern over data reliability (Congressional Research Service, 2025). Japan: zero response rates published, zero corrections applied -- institutional position that everything is functioning normally.

4.4 The Compulsory-Voluntary Divide

Legal compulsion is the strongest single predictor of survey resilience. The three voluntary-participation countries (US, UK, Canada) show 18 to 58 percentage-point response rate declines. The six compulsory-participation countries (France, Spain, Italy, Germany, Japan, Australia) maintain higher rates. But compulsion does not prevent decline: Italy's non-response increased 60 per cent in four years despite legal obligation, and Germany's rate reached 46.8 per cent during the 2020 overhaul. Only Australia maintains response rates above 90 per cent.

4.5 Measurement Obsolescence as the Shadow of Access

The relationship between the Measurement Obsolescence Hypothesis and the Access-Displacement Framework is not parallel. It is endogenous. MOH is the measurement shadow of the access process itself.

The causal mechanism operates as follows. When technology compresses the access gap, enabling individuals to participate in economic activity outside traditional employer-employee relationships through AI tools, platform architectures, and digital infrastructure,

the resulting economic activity falls precisely into the categories that household survey infrastructure cannot see. The Current Population Survey was designed to measure a labour market of stable employers and identifiable employees. The Current Employment Statistics survey counts jobs reported by establishments through payroll records. Neither instrument was designed to measure a freelancer using generative AI to deliver consulting services from a home office, a platform worker splitting time across three gig applications, or a micro-entrepreneur selling AI-augmented design work through an online marketplace. As documented in Section 4.2, Bracha and Burke (2023) estimated that official employment statistics understate actual employment by 0.25 to 5.1 percentage points, and Abraham et al. (2021) documented systematic misclassification between IRS records and CPS data on self-employment. These are not incidental measurement errors. They are structural consequences of an access process that creates economic participants faster than the measurement apparatus can be redesigned to count them.

The CES benchmark revision trajectory documented in Section 4.2 -- escalating from historical norms to 3.4 times the historical average -- is consistent with this interpretation. An accelerating benchmark revision in the absence of recession suggests that the composition of economic activity is shifting toward forms the survey captures poorly -- precisely the shift the access gap compression predicts.

This reframes the entire displacement measurement problem. The measurement apparatus does not passively record displacement or its absence. It actively converts adaptation into displacement by design. When a bank teller is laid off and becomes a freelance financial consultant using AI tools, the CES loses a payroll job and the CPS may or may not capture the new activity depending on how the respondent interprets "employed." When a displaced factory worker joins a platform as a gig driver, the establishment survey registers a job loss while the household survey may classify the worker as employed, unemployed, or out of the labour force depending on hours, regularity, and the respondent's self-identification. The same economic event (a worker adapting to technological change) can be measured as displacement, as employment, or as labour force exit, depending on which instrument is consulted and how the worker describes the transition.

The nine-country MOH evidence confirms that this is not a US-specific phenomenon. Every country examined shows the same structural pattern: declining survey reliability concentrated in precisely the populations most affected by technology-driven access: young adults, foreign-born workers, gig and platform participants, and those in non-standard housing. The measurement infrastructure is degrading fastest where the access process is most active, and the correlation is causal: the access process creates the populations that the measurement apparatus cannot see.

The implication for the displacement debate follows directly. Displacement predictions that rely on aggregate employment statistics, the very statistics whose measurement infrastructure is degrading, are building on a foundation that is eroding beneath them. The firm-level evidence for $G \sim A \sim G \sim D \sim$ independence (Adachi et al., 2024; Dauth et al., 2021;

Yang, 2022; Jeong and Jo, 2025) is sourced from corporate data and administrative records that are not affected by household survey degradation. The aggregate context (record-low unemployment, growing employment) relies on surveys that are. The honest statement: the firm-level evidence for independence is robust; the aggregate displacement narrative is built on data whose reliability is declining in proportion to the speed of the access process it claims to measure.

5. The Economic Forcing Function: Empirical Evidence

5.1 From Theory to Measurement

Henjoto (2026, Section 4.2) proposed the Economic Forcing Function as a deductive argument: the death of Moore's Law and Dennard scaling creates an escalating cost ceiling on centralised computation, forcing deployment toward hybrid hub-and-spoke architectures. This section provides the empirical data that the theoretical argument required.

5.2 The Training-Inference Cost Divergence

The Economic Forcing Function is quantifiable through two diverging trajectories. Frontier model training costs are growing at approximately 2.4 times per year (90 per cent confidence interval: 2.0 to 2.9 times, amortised hardware method; Cottier, Rahman, Fattorini, Maslej, Besiroglu, and Owen, 2024, arXiv:2405.21015). GPT-4 cost approximately \$78 million in cloud compute and approximately \$800 million in hardware acquisition. Billion-dollar training runs are projected by 2027.

Simultaneously, inference costs for equivalent capability are declining rapidly. The Stanford Institute for Human-Centered Artificial Intelligence AI Index (2025), drawing on Epoch AI data, reports that GPT-3.5-level performance (64.8 per cent on MMLU) cost approximately \$20 per million tokens in November 2022 and approximately \$0.07 per million tokens by October 2024, a 280-fold decline in two years. Epoch AI (2025) estimates a median inference cost decline of 50 times per year across benchmarks, accelerating to 200 times per year for trends beginning after January 2024.

The hardware substrate is decelerating. Hennessy and Patterson (2019), in their ACM Turing Award Lecture, documented that single-core performance improvement has fallen to approximately 3 per cent annually in the post-Dennard era, compared with approximately 50 per cent annually during the 1986 to 2002 golden age. GPU price-performance (floating-point operations per second per dollar) doubles approximately every 2.5 years for machine-learning-focused GPUs (0.14 orders of magnitude per year; Hobbhahn and Besiroglu, 2022, Epoch AI), three to four times slower than the Moore's Law era improvement rate of approximately 0.5 orders of magnitude per year. The A100-to-H100 transition delivered approximately three times the performance per dollar; subsequent generations show more modest compute-per-dollar gains alongside significant memory and bandwidth improvements. Leiserson, Thompson, Emer, Kuszmaul,

Lampson, Sanchez, and Schardl (2020), published in *Science*, established that post-Moore's performance gains must come from software, algorithms, and hardware specialisation, and that these gains will be "opportunistic, uneven, and sporadic."

This divergence (training costs rising exponentially while inference costs for equivalent capability collapse, against a decelerating hardware substrate) IS the Economic Forcing Function. It makes centralised frontier training increasingly expensive and distributed edge inference increasingly cheap. The rational economic response is the hub-and-spoke architecture described in Henjoto (2026, Section 5.2): centralised facilities for heavy training, miniaturised models deployed to edge devices for inference, with an integration layer connecting them.

The DeepSeek V3 case illustrates how the forcing function operates in practice. DeepSeek claimed \$5.58 million in training compute for its V3 model, but this figure covers only the final pre-training run and excludes research, ablation experiments, and infrastructure costs (SemiAnalysis (2025) estimated total programme cost at \$500 million to \$1.3 billion). The genuine efficiency innovations (mixture-of-experts architecture, 85 per cent GPU utilisation, custom CUDA kernels) are themselves a response to the cost ceiling: operating under US export controls that restrict access to frontier GPUs, DeepSeek was forced to extract maximum capability from constrained hardware. The forcing function produces efficiency innovation, not deployment failure.

5.3 The Miniaturisation Convergence

Every major artificial intelligence laboratory has converged on developing small models for edge deployment: Meta (Llama 3.2, 1 billion and 3 billion parameters), Google (Gemma 3, from 270 million to 4 billion parameters; Gemma 3n, approximately 4 billion effective parameters on mobile devices), Microsoft (Phi-4 mini, 3.8 billion parameters; Phi-4, 14 billion parameters), Hugging Face (SmolLM2, 135 million to 1.7 billion), Alibaba (Qwen2.5, 500 million to 3 billion), and Apple (approximately 3 billion parameters, compressed to 2 bits per weight). When all major participants converge on the same architectural response to a cost constraint, the constraint is real and the response is economically rational.

The performance of these small models illustrates the capability-cost decoupling. Phi-4 at 14 billion parameters achieves 84.8 on the MMLU benchmark and 80.4 on MATH, beating GPT-4o on mathematics while running locally. Google's Gemma 3 at 4 billion parameters outperforms its predecessor Gemma 2 at 27 billion parameters across benchmarks. Alibaba's Qwen2.5 at 500 million parameters outperforms Google's Gemma 2 at 2.6 billion parameters on mathematics and coding tasks -- a model one-fifth the size surpassing its larger predecessor. Apple's approximately 3-billion-parameter model, compressed to 2 bits per weight through quantisation-aware training, outperforms full-precision models two to three times its size on human evaluation.

Quantisation advances have been critical to the miniaturisation convergence. Apple pioneered 2-bit quantisation-aware training with learnable weight clipping, recovering quality lost through compression using low-rank adapters. Microsoft's BitNet b1.58

achieved 1.58-bit ternary quantisation (2 billion parameters, trained on 4 trillion tokens) with performance within one to two points of full-precision models at one-fifth the memory, and up to 6.25 times speedup on standard CPUs. MediaTek's Dimensity 9500 chipset integrates native BitNet 1.58-bit processing, doubling on-device large language model output speed for 3-billion-parameter models. Four-bit quantisation has become standard practice for deployment across multiple open-source formats, retaining approximately 90 to 95 per cent of model quality.

Algorithmic efficiency (the compute required for a given performance level) is halving approximately every eight months (Ho et al., 2024; arXiv:2403.05812), outpacing Moore's Law even in its heyday. Andreessen Horowitz estimates that the cost of equivalent-capability inference is declining approximately tenfold per year (Appenzeller, 2024). The edge AI market reached \$24.91 billion in 2025, projected to grow at 21.7 per cent compound annual growth rate to \$118.69 billion by 2033 (Grand View Research, 2025). Inference workloads consume over 55 per cent of AI-optimised infrastructure spending in early 2026, projected to reach 70 to 80 per cent by year-end (Deloitte, 2026). Gartner (2025a) predicts that by 2027, organisations will use small, task-specific AI models three times more than general-purpose large language models. Samsung targets 800 million Galaxy AI-enabled devices by end of 2026, quadrupling from 2024.

5.4 The Hyperscaler Capital Expenditure Trajectory

The financial stakes of the Economic Forcing Function are visible in hyperscaler capital expenditure. Combined spending by the eight largest technology companies grew from approximately \$172 billion in 2022 to approximately \$256 billion in 2024, with 2025 actual spending reaching approximately \$427 billion (IEEE Communications Society, December 2025, citing MUFG Americas). Approximately 75 per cent of 2026 capital expenditure is allocated to artificial intelligence infrastructure. Goldman Sachs projects cumulative hyperscaler capital expenditure of \$1.15 trillion from 2025 to 2027.

Cahn (2024), in Sequoia Capital's widely cited analysis, originally identified a \$200 billion revenue gap (September 2023), updated to \$600 billion (June 2024) as the gap widened rather than narrowed. Enterprise generative AI revenue reached approximately \$37 billion in 2025 (Menlo Ventures, 2025) against approximately \$427 billion in infrastructure spending -- an 11 to 14 times gap. JPMorgan estimates that \$650 billion in annual revenue would be required to deliver a 10 per cent return on projected cumulative AI investment of \$5 trillion through 2030 -- equivalent to \$34.72 per month from every iPhone user, or \$180 per month from every Netflix subscriber, in perpetuity.

The depreciation dynamics reveal the financial pressure. All major hyperscalers extended GPU useful life assumptions to reduce reported costs: Google from four to six years (saving \$3.9 billion in 2023), Microsoft to approximately six years, Meta from three to approximately 5.5 years through incremental extensions, and Oracle from four to six years. The collective impact was approximately \$18 billion per year in reduced depreciation charges, an accounting response to the cost ceiling. Amazon reversed this trend in

February 2025, shortening its GPU useful life back from six to five years and taking a \$700 million operating income charge plus \$920 million in accelerated depreciation for early equipment retirements, an explicit acknowledgment that the aggressive assumptions had been unsustainable. Nvidia's CFO responded by claiming that A100 GPUs shipped six years earlier were "still running at full utilisation," a self-serving assertion (Nvidia's business model depends on continued GPU demand) that nonetheless illustrates the infrastructure durability premise of the Economic Forcing Function.

This capex trajectory is approaching the limits of internal financing. UBS (2026) projects that hyperscaler capital expenditure in 2026 will consume nearly 100 per cent of operating cash flows (approximately \$577 billion across the five largest hyperscalers), compared with a ten-year average of approximately 40 per cent. Free cash flow (operating cash flow minus capital expenditure) has declined to approximately \$200 billion and is falling.^[^1]

5.5 Cross-Country Compute Infrastructure

The Economic Forcing Function operates unevenly across geographies. The United States hosts approximately 74.5 per cent of known GPU cluster performance (Epoch AI, 2025), approximately 54 per cent of the world's 1,189 hyperscale data centres (Synergy Research, Q1 2025), and approximately 53.7 gigawatts of installed data centre capacity (44 per cent of the global 122.2 gigawatts). China hosts approximately 14.1 per cent of GPU cluster performance, a gap widened by successive rounds of US export controls (October 2022, October 2023, December 2024, January 2025) restricting Nvidia's advanced GPUs.

The European Union's installed capacity of approximately 11.9 gigawatts represents a 4.5 to 1 ratio relative to the US. The EU has responded with sovereign AI initiatives (InvestAI targeting EUR 200 billion; Mistral AI; OpenEuroLLM for all 24 EU languages) but faces structural disadvantages in power availability, regulatory burden, and venture capital access.

Emerging compute hubs are developing rapidly. Japan allocated JPY 10 trillion (approximately \$65 billion) for semiconductors and AI through fiscal 2030. Saudi Arabia secured a \$5 billion DataVolt deal at NEOM. The UAE committed over \$15 billion through Microsoft and G42 partnerships. Singapore invested more than S\$1 billion over five years, including S\$500 million for GPUs.

The geographic concentration of compute infrastructure has implications for the Access-Displacement Framework. The Economic Forcing Function drives deployment toward hybrid architectures globally, but the centralised training component of the hub-and-spoke system is overwhelmingly concentrated in the United States and, to a lesser extent, China. Countries that lack sovereign compute capacity depend on US infrastructure for the training component, creating a technological dependency that may constrain their ability to develop nationally appropriate automation-employment policies.

Global data centre power consumption reached 415 terawatt-hours in 2024 (approximately 1.5 per cent of global electricity), projected to reach 945 terawatt-hours by

2030, equivalent to Japan's total electricity consumption (International Energy Agency, 2025). The energy dimension of the Economic Forcing Function reinforces the hybrid architecture: on-device inference consumes orders of magnitude less energy per query than data centre processing, creating environmental as well as economic incentives for distributed deployment.

The Economic Forcing Function creates deployment pressure; the question is how quickly institutions convert that pressure into workforce adjustment. This is the domain of the Organisational Absorption Rate.

6. Organisational Absorption Rate: From Mechanism to Rate

6.1 From Single Case to Cross-Country Framework

Henjoto (2026, Section 4.1) introduced the Organisational Absorption Rate using JPMorgan Chase as an upper-bound case: even with maximum CEO commitment, unlimited capital, and explicit willingness to restructure, the bank produced redeployment rather than displacement. Section 2.5 above identified five absorption mechanisms through which economies manage the automation-employment transition. This section moves from mechanism to rate: not how economies absorb automation, but how fast they absorb it, and what determines the speed.

6.2 The Rate Concept

The displacement debate implicitly assumes that technology capability translates into workforce displacement at the speed of deployment. If a bank can automate 30 per cent of its customer service queries, 30 per cent of customer service agents are displaced. The Access-Displacement evidence demonstrates that this assumption is wrong. Between capability arrival and workforce adjustment lies an institutional transmission mechanism with its own speed constraints. The Organisational Absorption Rate is the speed at which an economy converts automation capability into workforce adjustment (retraining, redeployment, managed attrition, or involuntary layoff) without operational degradation.

Three properties of the OAR emerge from the ten-country evidence, consistent with the broader technology diffusion literature (Comin and Hobijn, 2010) and the role of absorptive capacity in determining institutional readiness for new technologies (Cohen and Levinthal, 1990). First, the OAR is always slower than the technology capability frontier. AI capability is doubling approximately every six to eight months on key benchmarks (Ho et al., 2024). No institutional mechanism operates at that speed. Second, the OAR is institution-specific: it varies across the five absorption archetypes identified in Section 2.5, and the binding variable differs in each case. Third, the OAR determines the displacement timeline independently of the technology timeline. Displacement crisis narratives conflate these two timelines. The technology timeline is fast and accelerating. The absorption timeline is slow and institutionally determined. The gap between them is the window within which adjustment occurs, and every economy examined in the programme has filled that window

with some combination of demographic replacement, institutional negotiation, market reallocation, state administration, or military pipeline.

6.3 Binding Variables: What Determines Absorption Speed?

The cross-country evidence identifies five binding variables that govern absorption speed, one corresponding to each archetype.

Demographic replacement rate (Japan, South Korea, Taiwan). In demographic absorption economies, the binding variable is the rate at which retiring workers create vacancies that automation fills. Japan's working-age population is declining by approximately 500,000 to 600,000 per year. The mandatory retirement system (age 60 or 65, with re-employment at reduced wages) creates a predictable absorption channel: MUFG's 9,500 "automated" positions were absorbed through natural attrition over approximately five to seven years, at a rate determined by the retirement schedule of bubble-era hires (born 1960 to 1968, reaching mandatory retirement between 2023 and 2033). South Korea's working-age population began declining in 2020. Taiwan's total fertility rate fell below 0.8 in 2025. In all three economies, the demographic replacement rate exceeds the automation displacement rate, making the OAR question moot: there are not enough workers to displace.

Works council negotiation speed (Germany). In institutional absorption economies, what limits absorption speed is the pace of formal negotiation processes. A German Sozialplan (social compensation plan) negotiation between a works council and employer typically takes six to eighteen months. During this period, the terms of displacement are converted: involuntary layoffs become voluntary early retirement, retraining placements, or internal transfers. Kurzarbeit (short-time work) can be activated within weeks but addresses demand shocks, not structural automation. The Mittelstand apprenticeship system produces approximately 500,000 new apprentices per year, but retraining existing workers takes twelve to twenty-four months per cohort. Gathmann, Grimm, and Winkler (2024) found that AI is changing the task content of German occupations but that the institutional machinery is converting this into task reallocation rather than displacement. The binding constraint is negotiation throughput: how many Sozialplan processes the works council system can conduct simultaneously. As documented in Section 2.6, works council coverage has declined to approximately 42 per cent of West German and 37 per cent of East German employees (Kohaut and Schnabel, 2024), meaning the institutional absorption capacity is itself eroding, widening the gap between technology speed and absorption speed for the unprotected periphery.

Tripartite coordination speed (Singapore). Singapore's supply-side absorption model operates through the SkillsFuture system, which provides S\$500 in training credits per citizen aged 25 and over, supplemented by sectoral reskilling programmes coordinated through the tripartite National Wages Council. The bottleneck is programme throughput: how many workers can be reskilled per year and how quickly the training content adapts to new technology requirements. Singapore's National Service obligation (24 months for

males) delays labour market entry, creating an additional absorption buffer by limiting the supply of new workers competing with automated systems. Because the state can coordinate across employers, unions, and training providers simultaneously rather than negotiating firm by firm, Singapore's OAR is faster than in European institutional absorption economies.

Military pipeline intake (Israel). Israel's technology pipeline operates through compulsory military service, with elite units (Unit 8200 for signals intelligence, Talpiot for science and technology) recruiting approximately 1,000 to 1,500 graduates per year into technology roles. These graduates enter the civilian labour market with immediately deployable technical skills, reducing the training component of the absorption rate to near zero for the tech sector. What limits this mechanism is pipeline throughput: the annual cohort size. At approximately 1,000 to 1,500 per year, the pipeline can sustain the tech sector's growth (approximately 403,000 workers) but cannot expand rapidly. The limitation is that this mechanism serves only the tech-participating population (approximately 15 per cent): the Haredi and Arab-Israeli populations, comprising over 35 per cent of the total and growing faster than the secular Jewish population, are not connected to this pipeline.

State administrative capacity (China). China's OAR is bounded by the speed of state administrative intervention: how quickly the government can redirect employment through hukou policy adjustments, SOE hiring mandates, banking licence revocations (350 per year), and directed investment programmes. The Made in China 2025 programme demonstrates both the speed and the limitations: the state can accelerate automation faster than any market mechanism, but the administrative displacement management operates through the same bureaucratic channels that constrain all Chinese policy implementation: provincial compliance lags, information asymmetries between central and local government, and the tension between employment stability targets and productivity targets.

Market reallocation speed (Switzerland). Switzerland's OAR is bounded by labour market fluidity: the speed at which workers move from declining to growing sectors. Low employment protection (OECD EPL 2.10) removes institutional friction, allowing reallocation within months rather than years. But the mechanism depends partly on geographic displacement export: 403,000 cross-border workers absorb displacement pressure that would otherwise appear in Swiss unemployment statistics. The limiting factor is the absorptive capacity of the border regions, a finite resource that cannot scale indefinitely.

Table 5. Organisational Absorption Rate: Cross-Country Taxonomy

Archetype	Binding Variable	Estimated Timeline	Key Evidence
Demographic (JP, KR, TW)	Retirement replacement rate	5-10 years per cohort	MUFG 9,500 positions absorbed via bubble-era retirements (2023-2033)
Institutional (DE, SG)	Negotiation throughput /	6-24 months per Sozialplan	Works council coverage declining (80% to 49%);

	programme capacity		SkillsFuture throughput
Flexible (CH)	Market reallocation speed	Months (low EPL friction)	403,000 cross-border workers absorb displacement pressure
Segmented (IL)	Military pipeline intake	~1,000-1,500 per annual cohort	Unit 8200 pipeline serves only ~15% of population
State-mediated (CN)	Administrative intervention speed	Variable (state-directed)	Made in China 2025; 350 banking licences cancelled in 2025 (up from 198 in 2024)

6.4 The Klarna Test

The Klarna case (Section 3.4) provides an empirical test of what happens when an organisation attempts to exceed its OAR. Klarna reduced customer service headcount from approximately 3,000 to approximately 2,000 agents, celebrated the cost savings, and then reversed course when quality collapsed. It corroborates the OAR's central proposition: organisations that attempt to convert AI capability into workforce displacement faster than their internal complexity allows will experience operational degradation and reversal.

The Klarna reversal is not an isolated incident. Enterprise data consistently shows that organisational AI adoption is far slower than capability would permit, with high abandonment and reversal rates. Only 2 per cent of organisations have deployed agentic AI at scale, with 61 per cent remaining in exploration (McKinsey, 2025). Forty per cent of agentic AI projects are projected to be cancelled by 2027 (Gartner, 2025b). Only one in five companies planning agentic AI has mature governance frameworks in place (Deloitte, 2026, surveying 3,235 respondents across 24 countries). Over half of layoffs attributed to AI are estimated to be "quietly reversed" as organisations discover operational challenges (Forrester, 2026). This collective pattern describes an absorption rate governed by the OAR mechanism.

6.5 The Speed Gap: Technology Capability vs Institutional Absorption

The central insight of the OAR framework, extended to cross-country evidence, is that a structural speed gap exists between technology capability and institutional absorption in every economy examined. AI capability benchmarks are improving at approximately 2 to 4 times per year on standardised measures (Ho et al., 2024; Stanford Institute for Human-Centered Artificial Intelligence, 2025). Institutional absorption operates on fundamentally different timescales: demographic replacement cycles of five to fifteen years, works council negotiations of six to eighteen months, retraining programmes of twelve to twenty-four months, military pipeline cohorts of annual cycles.

This speed gap is not new. Every prior automation technology (mechanisation, electrification, computerisation) produced a similar gap between capability arrival and

institutional absorption, and in every case the institutional absorption rate, not the technology capability rate, determined the displacement outcome (Autor, 2015; Mokyr, Vickers, and Ziebarth, 2015; Bessen, 2019). The question is whether AI is different: whether it will overwhelm institutional absorption mechanisms that have successfully managed every prior technology wave.

The cross-country evidence permits a rough ordering of absorption speeds. Flexible absorption (Switzerland) operates fastest: low employment protection allows sectoral reallocation within months, though this speed is partly dependent on displacement export through cross-border workers. Institutional absorption (Singapore) operates at intermediate speed: tripartite coordination and centralised SkillsFuture throughput allow state-coordinated reskilling within one to two years per cohort, faster than Germany's firm-by-firm negotiation but bounded by programme capacity. Institutional absorption (Germany) is slower: Sozialplan negotiations take six to eighteen months per intervention, and each intervention covers one firm at a time. Demographic absorption (Japan, South Korea, Taiwan) operates on the slowest timescale (five to fifteen years per retirement cohort) but paradoxically faces the least urgency, because the retirement rate exceeds the automation displacement rate, making the speed question moot. State-mediated absorption (China) is variable: the state can impose automation faster than any market mechanism (Made in China 2025) but the administrative displacement management operates through bureaucratic channels with their own speed constraints. Segmented absorption (Israel) operates at the speed of the military pipeline (approximately 1,000 to 1,500 graduates per annual cohort) but serves only the tech-participating 15 per cent of the population.

None of these absorption speeds approaches the rate of AI capability improvement. But none needs to. The displacement gap is governed by the absorption rate, not the capability rate. As long as absorption mechanisms function (demographic replacement continues, institutional negotiation proceeds, markets reallocate, states administer, pipelines produce) the speed gap between capability and absorption is the window within which adjustment occurs, not the precipice over which displacement falls. The ten-country evidence suggests not yet. But the threshold identified by Jeon and Kwon (2024), where high automation levels produce significant employment decline even as low levels show no significant reduction, and the AI-specific dynamics documented in Israel (Debowy et al., 2025) indicate that the answer may change as AI adoption moves from the current 20 to 40 per cent range toward saturation.

7. Discussion

7.1 What the Evidence Means for the Access-Displacement Framework

The three empirical programmes reported here -- $G \sim A \sim G \sim D \sim$ independence across ten countries, international measurement obsolescence, and the quantified Economic Forcing Function -- strengthen each of the seven contributions advanced in Henjoto (2026). But

taken together, they produce findings that exceed what any programme demonstrates individually.

The $G\sim A\sim/G\sim D\sim$ independence claim is no longer US-centric. It is confirmed by six independent econometric studies across six countries (Adachi et al., 2024; Dauth et al., 2021; Yang, 2022; Jeong and Jo, 2025; Cali and Presidente, 2025; Giuntella et al., 2025), corroborated by firm-level data from Koch, Manuylov, and Smolka (2021) and Freeman et al. (2025), reinforced by cross-country robot-employment analysis (Graetz and Michaels, 2018; Gregory, Salomons, and Zierahn, 2022), and challenged only by China under conditions that do not generalise. The Measurement Obsolescence Hypothesis extends from a US-specific observation to a global structural finding: household surveys are degrading in all nine countries examined, through different mechanisms but toward the same outcome. The Economic Forcing Function is no longer a deductive argument from premises: the training-inference cost divergence, the miniaturisation convergence, and the capex trajectory are now quantified.

The interaction between the three programmes warrants particular attention. The Measurement Obsolescence Hypothesis is not a parallel observation to $G\sim A\sim/G\sim D\sim$ independence -- it is endogenous to the access process itself (Section 4.5). As technology compresses the access gap, the resulting economic activity falls into precisely the categories that household survey infrastructure cannot see. The measurement apparatus converts adaptation into displacement by design: a worker who leaves a payroll position for AI-augmented independent work registers as a job lost in the establishment survey and, if they do not respond to the household survey, as a worker who has left the labour force entirely. The nine-country evidence confirms that this is not a US-specific phenomenon: every country examined shows declining survey reliability concentrated in the populations most affected by technology-driven access. This means that displacement predictions calibrated against aggregate employment statistics are building on a foundation that is eroding in proportion to the speed of the access process they claim to measure. The firm-level evidence for $G\sim A\sim/G\sim D\sim$ independence is sourced from corporate data and administrative records unaffected by household survey degradation. The aggregate displacement narrative is built on data whose reliability is declining. These two observations are not in tension -- they are the same observation from different vantage points.

The Economic Forcing Function provides the supply-side explanation for why $G\sim A\sim/G\sim D\sim$ independence persists even as capability accelerates. Training costs rising at 2.4 times per year while inference costs for equivalent capability decline 280-fold in two years creates an economic architecture that structurally embeds human participation rather than eliminating it. The hub-and-spoke topology this produces -- centralised training, distributed edge inference, bidirectional integration -- is not one deployment option among many. It is the economically forced outcome of a cost structure that makes centralised replacement progressively less rational. Every major artificial intelligence laboratory converging simultaneously on sub-3 billion parameter edge models is the market confirming this analysis.

The US-outlier finding reframes the displacement debate at the level of the research programme itself. The automation-employment literature is overwhelmingly US-centric: Frey and Osborne (2017) use US occupational data; Acemoglu and Restrepo (2020) use US commuting zones; Goldman Sachs (2023) extrapolates global estimates from US-derived task analysis. Guarascio, Piccirillo, and Reljic (2025), in a meta-analysis spanning multiple countries, found that the United States is the only major economy where automation-employment studies consistently yield negative results. If the empirical base of the displacement tradition is drawn from the global outlier, then the tradition's predictions carry a systematic bias that no amount of methodological refinement within the US context can correct. The correction requires cross-country evidence of the kind this paper provides.

The cross-country findings are compatible with, and extend, the most rigorous recent work within the displacement tradition itself. Acemoglu (2024), in what is arguably the most careful quantitative treatment of AI's economic impact, found that only 4.6 per cent of all work tasks are cost-effectively automatable within a decade -- an order of magnitude below the headline "exposure" figures that institutional reports promote. This finding is consistent with the Organisational Absorption Rate mechanism documented here: cost-effectiveness, not capability, is the binding constraint, and cost-effectiveness is determined by the institutional and economic factors the OAR framework identifies. Acemoglu, Autor, and Johnson (2026), in the most comprehensive recent treatment of automation policy, proposed a five-category taxonomy of technological change -- distinguishing labour-augmenting, capital-augmenting, automating, expertise-levelling, and new-task-creating innovations -- and identified three market failures that bias artificial intelligence deployment toward automation rather than augmentation: insufficient consideration of task-level composition, inadequate investment in new task creation, and information failures that cause firms to over-automate. Their taxonomy operates at the technology level; the absorption archetype taxonomy developed here operates at the institutional level. The two are complementary: Acemoglu, Autor, and Johnson identify the supply-side distortions that bias technology toward displacement, while the five absorption mechanisms identify the demand-side institutional structures that determine whether displacement actually materialises. My reading of the cross-country evidence is that institutional structure, not technology type, is the stronger determinant of employment outcomes, though I recognise this interpretive claim exceeds what any single country study can demonstrate on its own.

7.2 Policy Implications

The finding that the United States is a global outlier in automation-employment dynamics carries direct policy implications. AI regulation and workforce policy are increasingly designed around US-centric evidence. The European Union's AI Act, the OECD's AI Principles, and the G7's Hiroshima AI Process all incorporate assumptions about displacement risk derived substantially from US research. If the US experience does not generalise, these policies may be miscalibrated for the majority of the world's economies.

The five absorption archetypes suggest that effective automation policy must be country-specific, not universal. The OECD (2023) reached a similar conclusion: artificial intelligence's labour market impact varies substantially across countries depending on institutional structures, workforce composition, and policy frameworks. Universal displacement predictions produce universal policy responses; the evidence suggests the problem is heterogeneous and the solutions must be too.

The archetype-specific implications are as follows.

Demographic absorption economies (Japan, South Korea, Taiwan). The central challenge here is not displacement but labour scarcity. Japan's working-age population is declining by approximately 500,000 to 600,000 per year; South Korea's began declining in 2020; Taiwan's fertility rate is the world's lowest. In these economies, the policy priority is not displacement prevention (there are not enough workers to displace) but managing the dual economy that automation concentration creates. Taiwan's Hsinchu corridor earns 65 per cent above the national median from 6.5 per cent of the workforce. South Korea's chaebols provide internal redeployment for core employees while platform workers absorb precarity (the Coupang pattern). The policy challenge is distributional: ensuring that automation's productivity gains reach beyond the demographic corridors where labour scarcity is most acute. Immigration policy becomes an automation policy lever: Israel's post-October 7 response demonstrated that even economies with extensive robotics capacity choose labour importation over automation when given the option. Japan, South Korea, and Taiwan face the same choice on expanding timescales.

Institutional absorption economies (Germany, Singapore). The rate-limiting factor is institutional throughput: how many workers can be retrained, redeployed, or negotiated into managed attrition per year. Germany's works council system converts displacement into Sozialplan negotiations that take six to eighteen months per intervention -- but that system now covers only 42 per cent of West German and 37 per cent of East German employees. Singapore's SkillsFuture system and tripartite coordination can reskill workers faster than Germany's firm-by-firm negotiation, but throughput is bounded by programme capacity and the speed at which training content adapts. The policy priority in these economies is maintaining and extending institutional absorption capacity as it erodes. Germany's declining collective bargaining coverage (80 per cent to 49 per cent over three decades) is not merely a labour market statistic; it is an erosion of the primary mechanism through which the economy has historically absorbed every prior automation wave. Policy that accelerates AI adoption without simultaneously strengthening institutional absorption capacity will produce displacement in the unprotected periphery: atypical workers, workers in firms without works councils, and workers in sectors where collective bargaining has already collapsed.

Flexible absorption economies (Switzerland). What governs absorption here is reallocation speed and the sustainability of displacement export. Switzerland's low employment protection (OECD EPL 2.10) allows rapid sectoral reallocation, and 403,000 cross-border workers absorb displacement pressure that would otherwise appear in Swiss

unemployment statistics. The policy implication is that Switzerland's favourable employment outcomes are partly dependent on the absorptive capacity of its border regions in France, Germany, and Italy -- a finite resource. If those economies face their own automation pressure simultaneously, Switzerland's displacement export channel narrows. Policy in flexible absorption economies must account for the geographic externality: low domestic unemployment may coexist with exported displacement that is invisible in national statistics.

Segmented dual economies (Israel). The core problem is the participation gap between the technology economy and the demographic growth economy. Israel's tech sector (403,000 workers, 11.5 per cent of the workforce) generates 17.3 per cent of GDP at a 2.8-times wage premium, but the fastest-growing populations, Haredi (13.9 per cent, projected 32 per cent by 2065) and Arab-Israeli (21.1 per cent), have tech participation rates of 3.5 per cent and 1.5 per cent respectively. The OECD projects a 12-percentage-point GDP per capita gap by 2065 between full integration and frozen participation rates. The policy priority is not automation regulation but bridging the segmentation: connecting the populations experiencing demographic growth to the economy experiencing technological growth. This is an education, cultural, and institutional challenge, not a technology challenge. The Taub Center's finding that 9 per cent of AI-adopting Israeli businesses report reducing their workforce suggests that as AI penetrates further into knowledge work, even the tech-participating population may face displacement pressure -- making the segmentation problem more urgent. Whether the 9 per cent figure represents the leading edge of a broader trend or a plateau remains to be seen.

Pre-automation economies (Indonesia, India). These economies face a different problem entirely: development-stage transition. Indonesia (robot density approximately 5 per 10,000) and India (approximately 7 per 10,000) are below the automation adoption threshold where the Access-Displacement Framework's $G \sim A \sim / G \sim D \sim$ dynamics engage. Displacement in these economies is driven by trade competition (Indonesia: Chinese import competition), oligarchic rent extraction, and structural informality (Indonesia's 59.4 per cent informal sector), not by technology. India's UPI/Jan Dhan revolution demonstrates that technology's primary economic role at this development stage is access creation for previously excluded populations, not substitution of existing workers. Policy designed around automation displacement -- the frameworks exported from US and EU research -- is misapplied in pre-automation economies. The appropriate policy priority is managing the development-stage transition: building human capital, formalising informal work, and ensuring that when automation adoption crosses the threshold identified by Cali and Presidente (2025), institutional absorption mechanisms exist to manage it.

Two cross-cutting policy implications apply across all archetypes.

First, measurement reform is a policy prerequisite. Section 4 demonstrated that household survey infrastructure is degrading in all nine countries examined. Policymakers cannot design effective automation policy with employment statistics that systematically miss the populations most affected by technological change. The Novel Temporal Labour

Indicators proposed in Henjoto (2026), measurement approaches that track economic transformation through behavioural, transactional, and linguistic signals rather than establishment headcounts, are not an academic suggestion. They are a practical necessity for any economy whose policy apparatus depends on employment data that is structurally failing.

Second, financial sector displacement misattribution (Section 3) has a specific policy implication: it systematically inflates perceived displacement risk, which in turn biases policy toward compensatory responses (income transfers, safety nets) rather than generative responses (new task creation, institutional absorption capacity, human-AI partnership investment). If 90 per cent of financial sector "AI displacement" decomposes into mergers, demographic attrition, and regulatory restructuring, then the displacement signal that policymakers use to calibrate urgency is nine-tenths noise. Policy calibrated to a misattributed signal will misallocate resources.

7.3 Limitations

These policy implications assume the evidence is robust. Honesty requires acknowledging what the evidence cannot yet establish.

The methodology is comparative-institutional, not quantitative-parametric. The programme draws on existing econometric studies but does not conduct new regressions. No panel data was constructed; no unified statistical test of $G \sim A \sim G \sim D \sim$ independence was performed across countries. The absorption archetype taxonomy is descriptive, not predictive: it classifies observed mechanisms but does not estimate their parameters.

The cross-country evidence is cross-sectional. Each country was examined at a point in time. Whether the $G \sim A \sim G \sim D \sim$ independence relationship holds over time -- particularly as AI adoption exceeds the current 20 to 40 per cent range -- remains empirically open. Jeon and Kwon (2024), studying Korean data, found non-linear effects: at low levels of automation, productivity effects offset displacement with no significant employment reduction, but at high levels, displacement dominates and employment declines significantly -- suggesting a threshold beyond which the relationship may reverse. Henjoto (2026, Section 6.4) identified this threshold effect as a critical boundary condition on the framework.

Three of the seven original contributions -- the Sampling Residualisation Hypothesis, the AGI-C versus AGI-R distinction, and the Hybrid Swarm Architecture -- have not yet been subjected to cross-country empirical validation. These remain for future work.

The programme examined ten countries. Important economies are absent: the United Kingdom (where LFS collapse prevents reliable employment measurement), Brazil, Mexico, and sub-Saharan African economies. The framework's applicability to these contexts is untested.

Israel's AI-first displacement dynamics (Section 2, Table 1) may represent an emerging pattern that the programme's robot-centric evidence base has not yet fully captured. The

Taub Center's finding that 9 per cent of AI-adopting Israeli businesses report reducing their workforce is an early signal that warrants longitudinal tracking.

7.4 Future Work

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

The threshold question. A critical boundary condition on the framework's central claim is whether $G \sim A \sim G \sim D \sim$ independence holds at higher adoption levels. Current evidence covers a 20 to 40 per cent adoption range (Crane, Green, and Soto, 2025). Jeon and Kwon (2024), studying Korean data, found non-linear effects: at low automation levels, productivity effects offset displacement with no significant employment reduction, but at high levels, displacement dominates and employment declines significantly. If a threshold exists beyond which institutional absorption mechanisms are overwhelmed, identifying that threshold empirically -- and determining whether it varies by absorption archetype -- is a high-priority research question for the framework.

The AI-robot divergence. Jeong and Jo (2025) found that AI increases employment by 4.8 per cent while robots decrease it by 3.3 per cent in the same Korean economy. This technology-type divergence -- if it replicates across countries -- would fundamentally alter the displacement debate, which has predominantly studied industrial robots. AI may operate through different mechanisms than robotics: augmenting cognitive tasks rather than substituting physical ones, creating demand for new task categories that robots do not, or interacting differently with the institutional absorption mechanisms identified here. Cross-country replication of the Jeong and Jo finding is a high-priority empirical test.

Distributional displacement. The distributional displacement pattern identified in Section 2.7 -- technology concentrating benefits in narrow corridors without aggregate job loss -- requires formal measurement. The framework currently defines $G \sim D \sim$ as an employment-count measure. The Taiwan and Israel evidence suggests that a distributional component is needed: automation can produce a dual economy in which the displacement gap remains closed while the welfare gap widens. Developing a formal distributional displacement index -- measuring the geographic and demographic concentration of automation's benefits alongside aggregate employment outcomes -- would extend the framework beyond its current formulation.

The population-scale hypothesis. China is the sole economy where worker-level $G \sim A \sim G \sim D \sim$ independence fails. The hypothesis that this failure requires three simultaneous conditions -- population exceeding one billion, unevenly distributed development, and rapid state-directed automation -- has $N=1$. India is the theoretical second candidate but is decades from China's automation level. Indonesia, the third large economy examined, is below the automation threshold entirely. Testing this hypothesis requires either longitudinal tracking of India's automation trajectory as it develops, or identification of additional large-scale economies that meet the conditions.

Longitudinal OAR tracking. Cross-institutional measurement of absorption rates from 2022 to 2028 would test whether the organisational absorption rate is accelerating, decelerating, or stable as AI tools mature. The prediction is stability or deceleration: firms investing more in AI but not displacing proportionally more workers, consistent with the Klarna pattern (Section 6.4) and the enterprise survey data showing increasing abandonment rates (S&P Global / 451 Research, 2025; Gartner, 2025b).

Measurement reform. Cross-country validation of the Sampling Residualisation Hypothesis -- particularly in the United Kingdom, where LFS collapse provides a natural experiment in measurement system failure -- would extend the measurement framework. The UK's decision to suspend and rebuild its labour force survey creates a before-and-after comparison that no other country can provide: if the rebuilt survey captures workers the old survey missed, the SRH gains direct empirical support.

Formal parametric testing. A unified panel data test of $G \sim A \sim G \sim D \sim$ independence across the ten countries examined here would move the framework from comparative-institutional to quantitative. The present paper's methodology is comparative; a panel estimation with country fixed effects, sector-level automation measures, and institutional absorption proxies would provide the statistical test that the framework's cross-generational evidence base -- limited to five to six technology transitions -- cannot support on its own.

An independent empirical contribution in a related domain is in preparation, extending the theoretical foundations established in Henjoto (2026).

8. Conclusion

For nearly a century, each major automation technology has produced predictions of mass displacement that did not materialise at the predicted scale or timeline (Mokyr, Vickers, and Ziebarth, 2015). This paper demonstrates that the pattern is not accidental. Across ten economies spanning four income levels, three continents, and five institutional configurations, technology-driven access and technology-driven displacement operate independently. The United States is the global outlier in automation-employment dynamics, not the representative case that the displacement literature assumes. Financial sector "AI displacement" -- the domain most favourable to technology-driven substitution -- is systematically misattributed in all ten countries examined. Household survey degradation is a structural global phenomenon that degrades the empirical foundation on which displacement predictions depend. And the Economic Forcing Function -- training costs rising exponentially while inference costs for equivalent capability collapse -- creates measurable pressure toward hybrid human-AI deployment architectures rather than centralised replacement.

Beyond validating the Access-Displacement Framework (Henjoto, 2026), this paper makes four contributions of its own. First, it identifies **distributional displacement** -- technology concentrating benefits in narrow corridors without aggregate job loss -- as a novel

extension of the framework that may explain the persistence of displacement anxiety amid favourable employment data. Second, it taxonomises **five absorption mechanisms** with distinct binding variables that determine absorption speed, establishing that the Organisational Absorption Rate is institution-specific and always slower than the technology capability frontier. Third, it establishes **immigration-automation substitution** as a universal pattern in high-income economies, where migrant labour systematically defers automation adoption. Fourth, it demonstrates that **measurement obsolescence is endogenous to the access process itself**: as technology compresses the access gap, the resulting economic activity falls into the blind spots of survey infrastructure designed for an earlier economy, causing the measurement apparatus to convert adaptation into displacement by design.

The framework does not predict that displacement is impossible. It predicts that displacement operates independently of access, that the binding constraints are institutional and demographic rather than technological, and that the current generation of displacement predictions -- built on US-centric evidence and calibrated to access-gap dynamics rather than displacement-gap dynamics -- will follow the same trajectory as their predecessors. Policy designed around US-centric displacement predictions is systematically miscalibrated for the majority of the world's economies. The evidence presented here supports this conclusion across ten economies.

Acknowledgments

The author acknowledges the use of artificial intelligence tools (large language models) in the preparation of this paper. These tools were used for cross-country evidence gathering, 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 judgments 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. DOI: 10.1086/712187.
- Acemoglu, D. (2024). The simple macroeconomics of AI. *NBER Working Paper* No. 32487.
- Acemoglu, D., Autor, D. and Johnson, S. (2026). Building pro-worker artificial intelligence. *NBER Working Paper* No. 34854. DOI: 10.3386/w34854.
- 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. and Restrepo, P. (2020). Robots and jobs: Evidence from US labor markets. *Journal of Political Economy*, 128(6), 2188-2244.

Acemoglu, D. and Restrepo, P. (2022). Demographics and automation. *Review of Economic Studies*, 89(1), 1-44. DOI: 10.1093/restud/rdab031.

Adachi, D., Kawaguchi, D. and Saito, Y. (2024). Robots and employment: Evidence from Japan, 1978-2017. *Journal of Labor Economics*, 42(2), 591-634. DOI: 10.1086/723205.

Anthropic (2025). Anthropic Economic Index: Uneven geographic and enterprise AI adoption. Anthropic Research. arXiv:2511.15080.

Appenzeller, G. (2024). Welcome to LLMflation -- LLM inference cost is going down fast. Andreessen Horowitz (a16z) Blog, November 12, 2024.

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. *Quarterly Journal of Economics*, 118(4), 1279-1333.

Beerli, A., Ruffner, J., Siegenthaler, M. and Peri, G. (2021). The abolition of immigration restrictions and the performance of firms and workers: Evidence from Switzerland. *American Economic Review*, 111(3), 976-1012. DOI: 10.1257/aer.20181779.

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. DOI: 10.1002/jae.3035.

Bessen, J.E. (2019). Automation and jobs: When technology boosts employment. *Economic Policy*, 34(100), 589-626. DOI: 10.1093/epolic/eiaa001.

Bracha, A. and Burke, M.A. (2023). Informal work and official employment statistics: What's missing? Federal Reserve Bank of Boston Working Paper No. 23-15.

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

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

Cali, M. and Presidente, G. (2025). Robots for economic development. *Labour Economics*, 96, 102731. DOI: 10.1016/j.labeco.2025.102731.

Capital Economics (2025). Taiwan economic data, reported via CNN, November 2025.

Chang, J.-H. and Huynh, P. (2016). ASEAN in transformation: The future of jobs at risk of automation. *ILO Bureau for Employers' Activities Working Paper* No. 9.

Cheng, Y.-H.A. and Loichinger, E. (2017). The future labor force of an aging Taiwan. *Population Research and Policy Review*, 36, 441-466. DOI: 10.1007/s11113-016-9423-z.

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. (2010). An exploration of technology diffusion. *American Economic Review*, 100(5), 2031-2059.

Congressional Research Service (2025). Current Employment Survey benchmark revisions. CRS In Focus IF12827.

Corlett, A. and Slaughter, H. (2024). Measuring up? Exploring data discrepancies in the Labour Force Survey. London: Resolution Foundation.

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

Dauth, W., Findeisen, S., Suedekum, J. and Woessner, N. (2021). The adjustment of labor markets to robots. *Journal of the European Economic Association*, 19(6), 3104-3153. DOI: 10.1093/jeea/jvab012.

de Leeuw, E.D., Hox, J.J. and Luiten, A. (2018). International nonresponse trends across countries and years: An analysis of 36 years of Labour Force Survey data. *Survey Methods: Insights from the Field*. DOI: 10.13094/SMIF-2018-00008.

Debowy, M., Winter, J., Epstein, G.S., Weiss, A. and Behar-Netanel, E. (2025). Employment trends and artificial intelligence in the Israeli labor market. Taub Center Policy Paper No. 02.2025.

Deloitte (2026). Technology, Media, and Telecommunications Predictions 2026. Deloitte Insights.

Deloitte Switzerland (2015). Man and machine: Robots on the rise? The impact of automation on the Swiss job market. Deloitte Switzerland, November 2015.

Dengler, K. and Matthes, B. (2018). The impacts of digital transformation on the labour market: Substitution potentials of occupations in Germany. *Technological Forecasting and Social Change*, 137, 304-316.

Devereux, Sir R. (2025). Independent review of the performance and culture of the Office for National Statistics. HM Government.

Epoch AI (2025). LLM inference prices have fallen rapidly but unequally across tasks. Epoch AI Data Insights.

Esping-Andersen, G. (1990). *The Three Worlds of Welfare Capitalism*. Princeton University Press.

Forrester (2026). Predictions 2026: The Future of Work. Forrester Research.

Freeman, R.B., Liu, X., Liu, Z., Song, R. and Xiong, R. (2025). The cause and consequence of robot adoption in China: Minimum wages and firms' responses. *Fundamental Research*, 5(4), 1759-1770. DOI: 10.1016/j.fmre.2022.07.016.

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

Gartner (2025a). Gartner predicts by 2027, organizations will use small, task-specific AI models three times more than general-purpose large language models. Gartner Press Release, April 9, 2025.

Gartner (2025b). Gartner predicts over 40% of agentic AI projects will be canceled by end of 2027. Gartner Press Release, June 25, 2025.

Gathmann, C., Grimm, F. and Winkler, E. (2024). AI, task changes in jobs, and worker reallocation. IZA Discussion Paper No. 17554.

Giuntella, O., Lu, Y. and Wang, T. (2025). How do workers adjust to robots? Evidence from China. *The Economic Journal*, 135(666), 637-652. DOI: 10.1093/ej/ueae086.

Goh, S.K., Wong, K.N., McNow, R. and Chen, L.-J. (2023). Long-run macroeconomic consequences of Taiwan's aging labor force. *Journal of Policy Modeling*, 45(1), 121-138. DOI: 10.1016/j.jpolmod.2023.01.006.

Goldin, C. and Katz, L.F. (2008). *The Race between Education and Technology*. Harvard University Press.

Goldman Sachs (2023). The potentially large effects of artificial intelligence on economic growth. Goldman Sachs Global Investment Research.

Goos, M. and Manning, A. (2007). Lousy and lovely jobs: The rising polarization of work in Britain. *The Review of Economics and Statistics*, 89(1), 118-133. DOI: 10.1162/rest.89.1.118.

Graetz, G. and Michaels, G. (2018). Robots at work. *Review of Economics and Statistics*, 100(5), 753-768.

Grand View Research (2025). Edge AI Market Size, Share & Trends Analysis Report. Grand View Research.

Gregory, T., Salomons, A. and Zierahn, U. (2022). Racing with or against the machine? Evidence on the role of trade in Europe. *Journal of the European Economic Association*, 20(2), 869-906.

Guarascio, D., Piccirillo, A. and Reljic, J. (2025). Robots vs. workers: Evidence from a meta-analysis. *Journal of Economic Surveys*, 39(5), 2254-2271. DOI: 10.1111/joes.12699.

Hall, P.A. and Soskice, D. (2001). *Varieties of Capitalism: The Institutional Foundations of Comparative Advantage*. Oxford University Press.

Han, J. and Oh, S. (2025). AI diffusion and youth employment. Bank of Korea Issue Note No. 2025-30.

Henjoto, V. (2026). Access without displacement: An Access-Displacement Framework for AI economic transformation. Zenodo. <https://doi.org/10.5281/zenodo.19051765>

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

Ho, A., Besiroglu, T., Erdil, E., Owen, D., Rahman, R., Guo, Z., Atkinson, D., Thompson, N. and Sevilla, J. (2024). Algorithmic progress in language models. arXiv:2403.05812. Presented at NeurIPS 2024.

Hobbhahn, M. and Besiroglu, T. (2022). Trends in GPU price-performance. Epoch AI.

IEEE Communications Society (2025). Hyperscaler capex \$600 billion in 2026. IEEE ComSoc Technology Blog, December 22, 2025.

IFR (2024). World Robotics 2024. International Federation of Robotics.

IMF (2025). Transforming the future: The impact of artificial intelligence in Korea. IMF Selected Issues Papers, 2025/013. DOI: 10.5089/9798229003193.018.

International Energy Agency (2025). Energy and AI. IEA, Paris.

Israel Democracy Institute (2024). 2024 Statistical Report on Ultra-Orthodox (Haredi) Society in Israel.

Israel Democracy Institute (2025). Haredim in Israel 2050: Demographic projections and economic and security scenarios.

Israel Innovation Authority (2025a). 2025 State of High-Tech Report.

Israel Innovation Authority (2025b). 2025 High-Tech Employment Status Report.

Jeon, Y. and Kwon, C. (2024). Impact of the automation by industrial robots on employment growth. *Journal of Social Science*, 63(3), 155-180. DOI: 10.22418/JSS.2024.12.63.3.155.

Jeong, J.H. and Jo, H.J. (2025). The effects of artificial intelligence and robotics on employment and wages in Korean manufacturing firms. *Weizenbaum Journal of the Digital Society*, 5(2). DOI: 10.34669/wi.wjds/5.2.5.

Kikuchi, S., Fujiwara, I. and Shirota, T. (2024). Automation and disappearing routine occupations in Japan. *Journal of the Japanese and International Economies*, 74, 101338. DOI: 10.1016/j.jjie.2024.101338.

Koch, M., Manuylov, I. and Smolka, M. (2021). Robots and firms. *The Economic Journal*, 131(638), 2553-2584. DOI: 10.1093/ej/ueab009.

Kohaut, S. and Schnabel, C. (2024). The demise of works councils in Germany. IZA Discussion Paper No. 17005.

Leduc, S., Oliveira, L.E. and Paulson, C. (2025). Do low survey response rates threaten data dependence? *Federal Reserve Bank of San Francisco Economic Letter*, 2025-07.

Lee, Y.S., Iizuka, T. and Eggleston, K. (2025). Robots and labor in nursing homes. *Labour Economics*, 92. NBER Working Paper No. 33116.

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. *Science*, 368(6495), eaam9744. DOI: 10.1126/science.aam9744.

McKinsey / QuantumBlack (2025). The state of AI in 2025: Agents, innovation, and transformation.

Menlo Ventures (2025). 2025 State of Generative AI in the Enterprise.

Mokyr, J., Vickers, C. and Ziebarth, N.L. (2015). The history of technological anxiety and the future of economic growth: Is this time different? *Journal of Economic Perspectives*, 29(3), 31-50.

MUFG Americas (2025). AI Chart Weekly: Financing the AI Supercycle. December 19, 2025.

Ni, B. and Obashi, A. (2021). Robotics technology and firm-level employment adjustment in Japan. *Japan and the World Economy*, 57, 101054. DOI: 10.1016/j.japwor.2021.101054.

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

OECD (2025). Long-term spending projections in Israel. OECD Publishing.

Oesch, D. and Piccitto, G. (2019). The polarization myth: Occupational upgrading in Germany, Spain, Sweden, and the UK, 1992-2015. *Work and Occupations*, 46(4), 441-469. DOI: 10.1177/0730888419860880.

Park, D. and Shin, K. (2025). Implications of artificial intelligence and robots for employment and labor productivity: Firm-level evidence from the Republic of Korea. ADB Economics Working Paper Series No. 769. DOI: 10.22617/WPS250038-2.

Piketty, T. (2014). *Capital in the Twenty-First Century*. Harvard University Press.

Rozelle, S. and Hell, N. (2020). *Invisible China: How the Urban-Rural Divide Threatens China's Rise*. University of Chicago Press.

S&P Global / 451 Research (2025). Voice of the Enterprise: AI & Machine Learning, Use Cases 2025. S&P Global Market Intelligence, October 2025.

Saxenian, A. and Hsu, J.-Y. (2001). The Silicon Valley-Hsinchu connection: Technical communities and industrial upgrading. *Industrial and Corporate Change*, 10(4), 893-920.

Schneider, F. (2022). New COVID-related results for estimating the shadow economy in the global economy in 2021 and 2022. *International Economics and Economic Policy*, 19, 299-313. DOI: 10.1007/s10368-022-00537-6.

SemiAnalysis (2025). DeepSeek debates: True training costs. SemiAnalysis Newsletter.

Spitz-Oener, A. (2006). Technical change, job tasks, and rising educational demands: Looking outside the wage structure. *Journal of Labor Economics*, 24(2), 235-270.

Stanford Institute for Human-Centered Artificial Intelligence (2025). AI Index Report 2025.

Swed, O. and Butler, J.S. (2015). Military capital in the Israeli hi-tech industry. *Armed Forces & Society*, 41(1), 123-141. DOI: 10.1177/0095327X13499562.

Synergy Research Group (2025). Hyperscale data centre count reaches 1,189 at end of Q1 2025.

Tsay, C.-L. and Lin, J.-P. (2001). Labor importation and unemployment of local workers in Taiwan. *Asian and Pacific Migration Journal*, 10(3-4), 505-534.

UBS (2026). UBS lifts forecast for big tech bond sales this year. Reported in Yahoo Finance, February 18, 2026.

Winters, J.A. (2011). *Oligarchy*. Cambridge University Press.

Wu, N. and Sun, Z. (2025). Little to lose: Exit options and attitudes towards automation in Chinese manufacturing. *The China Quarterly*, 262, 371-391. DOI: 10.1017/S0305741024001577.

Yang, C.-H. (2022). How artificial intelligence technology affects productivity and employment: Firm-level evidence from Taiwan. *Research Policy*, 51(6), 104536. DOI: 10.1016/j.respol.2022.104536.

Notes

^[companion] This is a companion paper to Henjoto (2026), "Access Without Displacement: An Access-Displacement Framework for AI Economic Transformation," available at <https://doi.org/10.5281/zenodo.19051765>. The present paper can be read independently but builds on the framework and definitions established there.

^[1] We note that Henjoto (2026) reported ">\$700 billion" in free cash flow for the five largest hyperscalers; this figure reflected operating cash flow, not free cash flow. The corrected free cash flow figure is approximately \$200 billion (2025), declining from approximately

\$237 billion (2024) as capital expenditure absorbs an increasing share of operating cash flow. Pivotal Research projects Alphabet's free cash flow alone declining 90 per cent in 2026. Barclays models Meta reaching negative free cash flow by 2027-2028. Hyperscalers issued over \$121 billion in bonds in 2025 alone, and for the first time hold more debt than cash. Amazon reversed its GPU depreciation schedule from six years back to five years in February 2025, taking a \$700 million operating income charge -- an explicit acknowledgment that the six-year assumption had been aggressive.