



3/2/2026

The Compression– Absorption Framework

A Stability Condition for Labor-
Compressing AI



iftikhar mahmud
INDEPENDENT SCHOLAR

Abstract

Artificial intelligence systems capable of automating decision-layer functions across sectors introduce a governance variable inadequately operationalized in contemporary AI policy: velocity. Institutional lag theory has long recognized that technology can outpace social adaptation (Ogburn, 1922; North, 1990; Pierson, 2000), and automation economics has documented task displacement dynamics (Autor, 2013; Acemoglu & Restrepo, 2020). What remains absent is a formal stability condition linking compression velocity to institutional absorptive capacity in a way that generates actionable governance thresholds. This paper provides that formalization. The principal macro-institutional risk of labor-compressing AI lies not merely in displacement, but in the rate at which compression may exceed institutional absorption — a distinction with significant policy implications. Modern states—regardless of ideology—rely on broad labor participation, distributed income flows, and identifiable authority structures for fiscal stability and political legitimacy. When technological deployment compresses labor demand faster than retraining systems, fiscal buffers, and political recalibration mechanisms can respond, a structural instability window emerges. The risk is not immediate collapse but gradual lock-in: permanent reorganization of production architectures, skill atrophy, capital concentration, and erosion of corrective optionality.

This paper introduces the Compression–Absorption Framework as a stability condition for labor-compressing AI. Stability requires that compression velocity (C) not exceed institutional absorptive capacity (A) over sustained periods. Persistent gaps generate cumulative lock-in pressure, increasing the cost of institutional correction. The paper advances a doctrine of Institutional Preemption: governance mechanisms for labor-compressing AI systems must be established prior to visible crisis. This approach neither predicts catastrophic unemployment nor opposes innovation. Rather, it reframes AI governance as equilibrium management in the face of acceleration asymmetry. If severe disruption does not materialize, anticipatory design may have contributed to that outcome. The central policy question is not whether technological transformation will occur, but whether institutional adaptation can keep pace with deployment velocity.

Keywords: Artificial Intelligence Governance, Labor-Compressing AI, Institutional Lag, Compression–Absorption Framework, Velocity Asymmetry, Path Dependency, Absorptive Capacity, Political Economy of Automation, Structural Lock-In, Precautionary Governance

Key Terms and Abbreviations

The following terms carry specific technical meanings in this paper. Readers from adjacent disciplines are encouraged to treat them as defined below rather than by their colloquial usage.

Compression velocity $C(t)$: A composite index measuring the rate at which labor-compressing AI systems displace, restructure, or reduce human labor inputs at time t . Components include task substitution rates, occupational entry bottleneck formation, and wage compression ratios, normalized to percentage-of-workforce equivalents per year.

Absorptive capacity $A(t)$: A composite index measuring institutional capacity to reintegrate displaced workers and recalibrate authority structures at time t . Components include retraining redeployment velocity, fiscal redistribution capacity, labor mobility rates, and administrative responsiveness.

Governance gap $G(t)$: The surplus of compression velocity over absorptive capacity, defined as $G(t) = \max\{0, C(t) - A(t)\}$. A positive $G(t)$ indicates that compression is proceeding faster than absorption at time t .

Cumulative lock-in pressure $P(T)$: The integral of $G(t)$ over a time horizon T , representing the accumulated structural strain that has not been absorbed. As $P(T)$ rises, the cost of institutional correction increases non-linearly.

Institutional preemption: A governance doctrine requiring that detection, modulation, and absorption mechanisms be established prior to visible crisis, rather than after structural lock-in has advanced.

Irreversibility window: The period during which $C(t) > A(t)$ is observable and $G(t)$ is rising, but $P(T)$ has not yet crossed the threshold at which low-cost institutional correction becomes politically or economically infeasible.

Labor-compressing AI: Artificial intelligence systems capable of automating decision-layer, analytical, and managerial functions across sectors simultaneously, in contrast to earlier automation waves that primarily addressed routine manual or clerical tasks.

High-autonomy systems: A descriptor for individual AI systems that integrate organizational authority, reduce human override capacity, and operate across multiple decision domains. Used as a technical property descriptor, not as a synonym for the category “labor-compressing AI.”

Abbreviations: CAF (Compression–Absorption Framework) | $C(t)$ (compression velocity) | $A(t)$ (absorptive capacity) | $G(t)$ (governance gap) | $P(T)$ (cumulative lock-in pressure) | DOSM (Department of Statistics Malaysia) | PERKESO (Social Security Organisation Malaysia) | HRD Corp (Human Resource Development Corporation Malaysia) | ASEAN (Association of Southeast Asian Nations) | ISCO (International Standard Classification of Occupations)

1. Introduction

Existing structural analyses of technological systems have identified architectural conditions under which systems drift from equilibrium toward extraction, asymmetry, and irreversibility—patterns documented across algorithmic platforms, monetary regimes, and industrial production models (Arthur, 1989; Pierson, 2000). These analyses identify *which* configurations generate systemic imbalance. The present paper addresses a distinct but equally foundational question: *under what temporal conditions* does deployment of such systems become destabilizing? Structural imbalance does not become a governance crisis instantaneously; it becomes one when deployment velocity outpaces institutional capacity to adapt. The Compression–Absorption Framework introduced here operationalizes this velocity dimension, enabling anticipatory governance before structural lock-in occurs.

The Contribution: From Institutional Lag to Operational Threshold

Technological transformation has historically restructured labor markets without dismantling the institutional foundations of modern states. Industrialization displaced manual labor while creating new manufacturing occupations; computerization automated clerical functions while expanding information sector employment; digital platforms reorganized commerce and communication (Mokyr, 1990; Goldin & Katz, 2008). Each wave generated labor market disruption, yet adaptation unfolded over timescales measured in decades rather than years. Workers transitioned between sectors through gradual cohort replacement and retraining; wage structures adjusted as new skill premiums emerged; educational systems expanded to meet shifting demand (Autor, 2013; Acemoglu & Restrepo, 2020). Fiscal institutions adapted through incremental policy reform—unemployment insurance, social safety nets, progressive taxation—calibrated to absorption capacity as displacement patterns became visible (Lindert, 2004). Political systems retained sufficient time to observe disruption, deliberate responses, and implement corrective mechanisms before structural lock-in occurred.

Advanced AI systems introduce a qualitatively different governance dynamic. Unlike prior waves that primarily automated routine tasks, these systems increasingly perform analytical, managerial, and coordination functions across diverse domains (Brynjolfsson & McAfee, 2014). More importantly, they scale at computational speed and can be deployed across sectors near-simultaneously, potentially compressing adaptation timelines that historically unfolded over decades into windows measured in years.

Public discourse tends to oscillate between catastrophic unemployment narratives and dismissals of technological pessimism. Both framings obscure the central structural issue. The principal governance variable introduced by advanced AI is velocity.

When deployment velocity persistently exceeds institutional adaptive capacity, a governance gap emerges — one that may harden into structural constraint before political systems recognize its cumulative effect.

Whether such acceleration occurs remains uncertain. Should deployment velocity increase—through broader adoption or emergence of more capable architectures—the compression-absorption gap could widen at rates that overwhelm institutional responses. This possibility, rather than any specific capability forecast, justifies building velocity-monitoring infrastructure and response mechanisms in advance. If severe compression does not

materialize, anticipatory design may have contributed to that outcome. If it does materialize, delayed institutional preparation becomes prohibitively costly once structural lock-in has occurred.

This paper does not claim to have discovered that technology can outpace institutions; Ogburn (1922) established that, and North (1990), Pierson (2000), and Blanchard and Summers (1986) elaborated it extensively. The contribution is narrower and more specific: it provides a governance-operable formalization of institutional lag as a triggerable threshold rather than a descriptive tendency. Institutional lag theory identifies that temporal asymmetry between technology and institutions exists. The Compression–Absorption Framework specifies when that asymmetry crosses a threshold warranting preemptive governance response, through what mechanism it generates lock-in, and what observable conditions (persistent $C(t) > A(t)$, rising $G(t)$) should activate policy. This generates two implications that institutional lag theory does not supply. First, a falsifiable activation criterion: if observable indicators do not show $C(t)$ persistently exceeding $A(t)$ following widespread labor-compressing AI deployment, the framework’s preemptive governance prescription is not warranted. Second, a governance instrument calibration logic: the threshold at which modulation mechanisms are triggered is defined by the relationship between $G(t)$ and jurisdiction-specific $A(t)$ capacity, not by ad-hoc political judgment. Whether this formalization constitutes theoretical extension or governance-level operationalization is a methodological question the paper does not seek to resolve; the claim is practical rather than meta-theoretical.

The argument does not depend on predictions of mass unemployment or specific capability forecasts. It rests on asymmetric risk: the cost of early institutional preparation is limited relative to the cost of delayed correction once structural lock-in has occurred. Whether technological acceleration continues, plateaus, or accelerates unpredictably, velocity-sensitive governance infrastructure provides institutional resilience across scenarios. In this sense, anticipatory governance functions as institutional insurance.

Temporal Scope and Data Status

A note on temporal scope is required. The Malaysian data points referenced in Section 9 (covering the period 2015–2025, with some 2025 indicators) reflect official government statistics and legislative developments available as of early 2026. Where 2025 data points are cited, they derive from published official releases — DOSM Labour Force Statistics (November 2025), DOSM Employee Wages Statistics (Q2 2025), MOF budget communications, and PERKESO media releases — and are treated as verifiable near-contemporaneous observations, not forward projections. The framework itself, however, is designed for anticipatory governance: its diagnostic logic must be operationalizable before crisis-level data exists. Policymakers applying the Compression–Absorption Framework in real time will necessarily be working with the most recent available indicators rather than complete historical series. This is a feature of the framework’s design, not a limitation of the present analysis. The illustrative application in Section 9 demonstrates precisely this forward-leaning use case — identifying the gap between available data and full operationalization as part of the research agenda.

Relationship to Existing AI Governance Frameworks

The Compression–Absorption Framework operates as a complement to, not a substitute for, existing AI governance instruments. It addresses a dimension that current frameworks do not systematically cover: the rate of institutional displacement relative to absorptive capacity. The EU AI Act (2024) classifies AI systems by risk category — unacceptable, high, limited, and minimal — with attendant conformity requirements. This is a static risk taxonomy applied to individual system characteristics. The Compression–Absorption Framework adds a dynamic, macro-level dimension: even systems that individually pass risk-tier assessments can, in aggregate and at scale, generate governance gaps if deployed at velocities that exceed institutional absorption. The OECD AI Principles (2019, updated 2023) emphasize transparency, accountability, robustness, and human-centred values — again, system-level criteria. The present framework operates at the sector and national level, asking whether the cumulative deployment pace across systems is consistent with institutional stability, irrespective of individual system compliance. National AI strategies — including Malaysia’s National AI Roadmap (2021–2025) and analogous ASEAN frameworks — address industrial policy, talent pipelines, and adoption targets, but do not provide velocity-sensitive governance mechanisms triggered by compression dynamics. The Compression–Absorption Framework is designed to interface with these existing instruments: the EU AI Act’s high-risk tier classifications can serve as inputs to $C(t)$ measurement; OECD accountability principles apply to the Human Agency Safeguard Protocol; national AI roadmap adoption targets provide deployment velocity data for $G(t)$ calculation. The relationship is supplementary and instrumentally compatible.

2. Literature Review

Institutional Lag, Path Dependency, and Labor Compression

The argument advanced here builds upon five established strands of scholarship: institutional lag theory, path dependency and increasing returns, task-based labor economics, political legitimacy theory, and precautionary governance.

2.1 Institutional Lag and Adaptive Capacity

The concept of institutional lag has deep roots in social theory. Ogburn (1922) first articulated "cultural lag"—the phenomenon whereby technological change outpaces social and institutional adaptation, creating periods of maladjustment until norms and structures catch up. North (1990) demonstrated that institutional evolution is fundamentally path dependent, constrained by historical trajectories and existing coordination mechanisms. Formal rules may change incrementally, but informal constraints—embedded practices, organizational routines, enforcement mechanisms—evolve far more slowly. Pierson (2000) extended increasing returns dynamics into political systems, showing how early institutional choices create self-reinforcing mechanisms that make subsequent change progressively more difficult. As institutions develop complementarities and constituencies organize around existing arrangements, the political cost of reform escalates over time.

These literatures converge on a core insight: institutions adapt more slowly than technological systems. The temporal asymmetry exists across contexts, but its magnitude varies with deployment speed. Advanced AI systems may amplify this lag by compressing deployment cycles that historically unfolded over decades into timeframes measured in years, potentially exceeding institutional recalibration capacity.

2.2 Path Dependency and Structural Lock-In

Arthur (1989) and David (1985) demonstrated how technological systems exhibit increasing returns to adoption. Early choices—even those driven by historical accident rather than efficiency—can become locked in as networks, complementary investments, and user learning create switching costs. The QWERTY keyboard persisted despite ergonomic alternatives because coordination costs and skill specificity made transition prohibitively expensive. Pierson (2000) showed that political institutions follow similar dynamics: once authority structures, fiscal arrangements, and distributional coalitions form around particular organizational architectures, reversal becomes costly and politically constrained. Institutional change faces collective action barriers, sunk cost resistance, and coordination failures that technological change does not.

Applied to advanced AI, large-scale reorganization around labor-light production systems may generate reinforcing feedback loops. Capital deepening increases returns to further automation; displaced workers experience skill atrophy that raises reintegration costs; political coalitions form around concentrated actors who benefit from existing arrangements (Korinek & Stiglitz, 2019). If deployment proceeds faster than corrective mechanisms can activate, the window for low-cost institutional adjustment may close before political systems recognize the cumulative effect. This is not technological determinism—it is recognition that timing matters for institutional reversibility.

2.3 Task-Based Automation and Labor Market Frictions

Task-based labor economics demonstrates that automation substitutes for specific tasks within occupations rather than eliminating entire occupations wholesale (Autor et al., 2003). This granular framework reveals that technological change reshapes the task composition of work rather than simply reducing total labor demand. Acemoglu and Restrepo (2019) formalize this dynamic through two countervailing forces: displacement effects (automation reducing labor demand for specific tasks) and reinstatement effects (creation of new tasks that restore labor demand). The net employment outcome depends on the relative strength of these forces. Where displacement outpaces reinstatement, labor share declines; where reinstatement dominates, new categories of work absorb displaced workers.

Empirical evidence on recent automation waves shows mixed results. Autor and Dorn (2013) documented wage polarization in U.S. labor markets between 1980-2005, with employment growth concentrated in high-skill and low-skill occupations while routine-task-intensive middle-skill jobs declined. Acemoglu and Restrepo (2020) examined industrial robot adoption in the United States and found significant negative effects on employment and wages in affected local labor markets, with limited evidence of offsetting job creation in the short to medium term. Each robot added per thousand workers reduced the employment-to-population ratio and wages, effects concentrated in manufacturing regions.

Blanchard and Summers (1986) introduced hysteresis theory in the context of European unemployment, demonstrating how temporary shocks can produce permanent effects through skill depreciation, insider-outsider dynamics, and labor force detachment. Subsequent research confirmed that workers displaced during recessions experience persistent wage penalties and reduced employment probability even after economic recovery (Davis & von Wachter, 2011; Kroft et al., 2013). Long-term unemployment generates scarring effects that compound over time.

These literatures establish that labor market adjustment to technological change is neither instantaneous nor frictionless. However, they analyze displacement magnitudes and long-run equilibria rather than adjustment dynamics. The present paper extends this foundation by shifting analytical focus from the scale of automation to the velocity of deployment relative to institutional adaptation capacity. If task displacement occurs gradually, reinstatement mechanisms and retraining systems operate within established timeframes. If displacement accelerates beyond historical rates—particularly across multiple sectors simultaneously—the friction between compression and absorption becomes the primary governance variable.

2.4 Political Legitimacy and Economic Inclusion

Lipset (1959) established that economic development and widespread prosperity correlate with democratic stability, though his framework emphasizes aggregate wealth rather than labor participation specifically. The connection between employment and political legitimacy operates through multiple channels: fiscal capacity depends on broad tax bases generated by wage income (Beramendi & Rueda, 2007); political voice correlates with economic participation (Schlozman et al., 2012); and institutional trust erodes when large populations become economically marginalized (Algan et al., 2017).

Rodrik (2011) identified structural tensions between economic integration, national sovereignty, and democratic governance. When economic outcomes decouple from domestic policy control—as occurs under certain globalization conditions—political legitimacy faces pressure. Labor-compressing AI may introduce analogous dynamics: if decision-layer authority migrates into computational systems controlled by concentrated actors, domestic policy instruments become constrained regardless of democratic preferences.

Fukuyama (2014) examined how institutional decay occurs when formal authority structures persist while operational power shifts elsewhere. His framework emphasizes that legitimacy requires not merely legal authority but effective institutional capacity to translate public preferences into outcomes. This directly informs the compression-absorption dynamic: if labor compression proceeds faster than institutions can recalibrate participation mechanisms, a legitimacy gap emerges independent of constitutional form.

These frameworks converge on a core proposition: modern state stability depends on maintaining alignment between economic participation structures and political authority. When technological transformation disrupts this alignment faster than institutions adapt, corrective capacity diminishes.

2.5 Asymmetric Risk and Precautionary Governance

While anticipatory governance frameworks provide general principles for emerging technologies (Guston, 2014; Marchant et al., 2008), the doctrine of Institutional Preemption advanced here is specifically calibrated to velocity-sensitive systems. The compression-absorption framework formalizes when anticipatory intervention becomes necessary: when deployment velocity threatens to exceed institutional adaptation capacity.

The precautionary principle and asymmetric risk frameworks (Taleb, 2012; Sunstein, 2005) justify early safeguards where downside risks are systemic and potentially irreversible. Financial regulation after 2008 illustrates this logic in practice. Capital buffer requirements, stress testing, and macroprudential oversight were instituted not because systemic collapse was deemed certain, but because recovery costs proved catastrophic once failure occurred (Admati & Hellwig, 2013). The regulatory response prioritized resilience over optimization.

Labor-compressing AI exhibits similar asymmetric risk characteristics. The cost of premature institutional preparation—velocity monitoring infrastructure, absorptive capacity buffers, governance frameworks—remains bounded and potentially reversible. Monitoring systems that prove unnecessary can be deactivated; fiscal reserves can be redirected; regulatory frameworks can be adjusted as technological trajectories become clearer (Sunstein, 2005). Conversely, the cost of delayed intervention escalates nonlinearly once compression dynamics exceed institutional capacity. Path dependency effects lock in production architectures (Pierson, 2000); skill atrophy raises reintegration costs; political coalitions form around concentrated actors (Korinek & Stiglitz, 2019). If deployment velocity compresses adaptation timelines from decades into years, institutional correction becomes progressively costlier as entrenchment deepens.

Under deep uncertainty about AI capability trajectories and deployment patterns, asymmetric risk logic favors early preparation over reactive response. The compression-absorption framework operationalizes this logic: institutional preemption activates when observed compression velocity approaches or exceeds measured absorption capacity, regardless of whether current systems represent incremental progress or precursors to more transformative capabilities.

3. The Compression–Absorption Framework Formalizing the Stability Condition

Technological change becomes destabilizing when its deployment velocity exceeds institutional adaptation capacity. To formalize this relationship, this paper introduces the Compression–Absorption Framework. A prior modeling-level clarification is necessary. $C(t)$ and $A(t)$ are not observed scalar quantities; they are theoretical constructs that must be operationalized as composite indices through explicit aggregation of heterogeneous components. The stability condition $C(t) \leq A(t)$ is therefore a structural diagnostic rather than a directly measurable equation. This places the framework at the level of intermediate theory: more precise than qualitative institutional lag accounts, less tractable than fully specified econometric models. The framework’s practical function is to define the measurement agenda — specifying what must be observed, at what granularity, and with what aggregation logic — prior to empirical calibration. Appendix A4 provides a worked numerical operationalization. The integral $P(T) = \int_0^T G(t)dt$ is approximating: without continuous real-time $G(t)$ measurement, $P(T)$ must be estimated from periodic snapshot data, introducing smoothing error. These are not incidental limitations; they define the framework’s current maturity level.

Let: $C(t)$ represent the rate of labor compression at time t , including displacement, task reduction, wage compression, and entry bottlenecks.

Let: $A(t)$ represent institutional absorptive capacity at time t , including retraining velocity, fiscal redistribution capacity, labor mobility, and administrative responsiveness.

Core Stability Condition $C(t) \leq A(t)$

When compression does not exceed absorption, reallocation can occur without structural drift.

The inequality compares compression velocity to absorption velocity; both are measured in consistent units (e.g., percentage of workforce per year). In practice, $C(t)$ may be constructed as a composite index normalized to percentage-of-workforce equivalents, enabling components with heterogeneous raw units—wage compression ratios, entry bottleneck rates, displacement counts—to be aggregated on a common scale prior to comparison with $A(t)$.

Define the governance gap: $G(t) = \max\{0, C(t) - A(t)\}$

If $G(t) > 0$, the instability window opens. The inequality $C(t) \leq A(t)$ is deliberately parsimonious. As a structural diagnostic, it abstracts away from heterogeneity in compression types, non-linearities in absorption, and the distributional concentration of displacement across workers and regions. These are substantive limitations: the composite index construction described below and in Appendix A4 is one approach to partial operationalization, but precise empirical calibration — particularly of the lock-in probability function relating $P(T)$ to irreversibility risk — awaits longitudinal data that do not yet exist at the required granularity. The framework’s analytic value lies in the precision of its diagnostic logic rather than the precision of its numerical outputs.

Cumulative Lock-In Pressure Over time horizon T : $P(T) = \int_0^T G(t) dt$

As cumulative pressure $P(T)$ increases, the probability of structural lock-in rises through increasing returns dynamics.

3.1 The Endogeneity of Absorptive Capacity

Importantly, $A(t)$ is endogenous. Innovation expands absorptive capacity through two channels. Market mechanisms generate new tasks and complementary labor demand—what Acemoglu and Restrepo (2019) term "reinstatement effects." General-purpose technologies create new industries and occupational categories that can absorb displaced workers (Brynjolfsson et al., 2021). Simultaneously, institutional mechanisms—retraining systems, fiscal stabilization, mobility infrastructure—enable workers to access emerging opportunities. The framework does not assume pessimistic outcomes; technological progress historically generates new forms of productive work alongside automation.

This endogeneity has a critical implication for the framework's governance logic. $A(t)$ is not merely responsive to market forces; it is also policy-responsive. When $G(t)$ becomes observable—when compression is seen to exceed absorption—deliberate institutional investment can increase $A(t)$ directly: scaling retraining pipelines, expanding fiscal buffers, accelerating mobility infrastructure. This feedback loop is the mechanism through which institutional preemption operates. The governance imperative is to act on $A(t)$ in response to rising $G(t)$ signals before $P(T)$ crosses entrenchment thresholds, not to treat $A(t)$ as a fixed parameter. The stability condition $C(t) \leq A(t)$ is therefore a governance target as much as a diagnostic: the framework prescribes raising $A(t)$ toward $C(t)$, not merely measuring the gap between them.

However, natural expansion proves structurally insufficient under velocity asymmetry. Three dynamics limit the stabilizing capacity of market-driven absorption:

Temporal lag asymmetry. Reinstatement operates over longer time horizons than displacement. Task-based research demonstrates that automation substitutes for existing tasks rapidly, while the creation and institutionalization of new task categories unfolds gradually (Autor, 2013; Acemoglu & Restrepo, 2018). Even where aggregate employment eventually balances, transition periods generate hysteresis effects—skill atrophy, labor force detachment, permanent wage scarring—that persist beyond initial shocks (Blanchard & Summers, 1986; Davis & von Wachter, 2011). Workers experiencing prolonged displacement face compounding barriers to reintegration. Markets may create new roles, but not before cumulative pressure $P(T)$ crosses entrenchment thresholds. The question is not whether reinstatement occurs, but whether it occurs fast enough to prevent structural lock-in during the transition.

Distributional mismatch. New tasks do not automatically route to those most compressed. Emerging roles may require different skills, concentrate in different regions, or favor different demographic profiles than displaced workers can readily access. Autor and Dorn (2013) documented wage polarization effects where job growth concentrated in high-skill and low-skill occupations while middle-skill workers faced persistent displacement. Geographic concentration compounds this dynamic: new industries cluster in specific regions while displaced workers face mobility constraints. Market-driven expansion thus generates aggregate balance without ensuring individual-level reallocation. Youth cohorts, mid-career

workers with obsolete specializations, and geographically immobile populations experience compression even when headline employment statistics appear stable.

Coordination barriers. Natural task creation does not include the complementary infrastructure required for efficient reallocation. Workers need certification systems to credential new skills, information networks to identify emerging opportunities, and mobility mechanisms to access geographically distant roles (Brynjolfsson et al., 2021). These complementary investments exhibit public goods characteristics and coordination externalities that markets undersupply. Institutional investment remains necessary even when markets generate new roles. Without coordinated retraining pipelines, mobility subsidies, and information infrastructure, workers cannot efficiently transition from compressed to expanding sectors.

Institutional preemption is thus necessary not because markets fail to expand $A(t)$, but because they expand it **too slowly** and **too unevenly** relative to $C(t)$ velocity under high-speed deployment conditions. The governance gap $G(t)$ emerges from acceleration asymmetry, not absolute capacity limits. Where deployment compresses decision-layer functions faster than natural reinstatement plus institutional facilitation can absorb, lock-in pressure accumulates. This is the central diagnostic insight: **the question is not how many jobs disappear, but whether compression velocity outpaces adaptive capacity.**

3.2 Manifestations of Labor Compression

Labor compression rarely manifests first as visible unemployment. More commonly, it unfolds through subtler dynamics: narrowing entry pathways, wage stagnation, thinning managerial layers, and task substitution within existing occupations. Task-based labor economics demonstrates that automation typically restructures occupations internally before eliminating them entirely (Autor, 2013; Acemoglu & Restrepo, 2018). These incremental adjustments may obscure deeper structural reorganization underway. Moreover, persistent labor market shocks can generate hysteresis effects, where temporary displacement becomes structurally embedded through skill erosion and labor detachment (Blanchard & Summers, 1986).

Labor-compressing AI systems differ from earlier automation waves in their capacity to compress decision-layer functions across sectors simultaneously. When firms integrate such systems into core managerial, analytical, and coordination processes, the transformation is not merely operational; it is architectural. Capital allocation shifts toward automation infrastructure, organizational hierarchies flatten or centralize, and strategic planning becomes increasingly dependent on algorithmic systems. As general-purpose technologies historically reshape production systems, their integration often produces organizational restructuring beyond immediate efficiency gains (Brynjolfsson et al., 2021).

3.3 Path Dependency and the Rising Correction Cost Window

Economic theory suggests that such architectural shifts exhibit path dependency and increasing returns dynamics. Early adoption generates learning effects, network advantages, and capital deepening that reinforce subsequent investment decisions (Arthur, 1989; David, 1985). Political institutions similarly display increasing returns, where early structural shifts constrain later reversal (Pierson, 2000). As firms invest in high-autonomy systems, investors align expectations around labor-light growth models, and skills associated with displaced

decision-layer roles atrophy. Over time, these reinforcing mechanisms raise the cost of re-expanding human-intensive structures.

The correction cost window describes the period during which corrective institutional intervention remains feasible at manageable cost — a period defined not by absolute irreversibility but by the political economy of reform under increasing structural constraint. When cumulative compression pressure—formalized earlier as $P(T)$ —remains limited, policy adjustment can recalibrate labor participation and authority distribution without systemic disruption. As this pressure accumulates, however, correction becomes increasingly disruptive, economically costly, and politically constrained. This dynamic is consistent with path-dependent institutional entrenchment, in which increasing returns reduce the feasibility of policy reversal (Pierson, 2000). Historical evidence does demonstrate that automation architectures can be partially reversed or restructured — antitrust action has broken up platform monopolies, industrial policy has re-onshored manufacturing, and financial regulation has curtailed automated trading after demonstrable systemic harm. Lock-in is therefore probabilistic rather than absolute: it describes a rising cost trajectory rather than categorical irreversibility. The correction cost window closes not because reversal becomes impossible, but because its political and economic cost rises to levels that democratic systems rarely sustain without precipitating crisis. The threshold between reversible and entrenched is empirically context-dependent — determined by $P(T)$ relative to jurisdiction-specific political tolerance, institutional flexibility, and fiscal capacity — and cannot be specified in advance without empirical calibration. The framework’s claim is the conditional one: if $P(T)$ is allowed to accumulate unchecked, the probability of low-cost correction declines non-linearly.

The core danger, therefore, is not unemployment alone. It is the erosion of corrective optionality. Corrective optionality refers to the capacity of institutions to rebalance labor participation, redistribute productivity gains, and reassert distributed authority once structural compression has advanced. When production systems reorganize around high-autonomy architectures, that optionality narrows in ways analogous to technological lock-in (Arthur, 1989).

3.4 Authority Redistribution and Constitutional Implications

This dynamic carries implications beyond labor economics. Decision-layer compression redistributes authority. Functions such as credit evaluation, regulatory enforcement, healthcare triage, welfare eligibility, and hiring increasingly rely on computational systems. Even when human oversight formally remains, operational discretion may become algorithmically mediated. Authority shifts from distributed human agents to centralized system operators. Institutional theory emphasizes that authority structures and accountability chains are foundational to governance coherence (North, 1990; Fukuyama, 2014).

Sovereignty in modern states depends not only on territorial control but on identifiable chains of authority and distributed participation in economic life. Economic participation has long been linked to political legitimacy (Lipset, 1959). When decision-layer authority migrates into opaque computational architectures at scale, governance becomes increasingly dependent on system designers and operators. AI governance thus intersects with constitutional design, raising questions not only of safety and bias but of authority concentration and institutional durability (Rodrik, 2011; Fukuyama, 2014).

The risk trajectory is gradual rather than catastrophic. Modern states rarely collapse due to technological change. More frequently, they experience entrenched wage stagnation, mobility decline, capital concentration, and institutional fatigue (Fukuyama, 2014). Structural lock-in is less visible than crisis, yet more enduring. Preventing such entrenchment requires attention to deployment velocity before the irreversibility window closes.

3.5 From Framework to Doctrine: The Case for Preemption

The stability condition $C(t) \leq A(t)$ formalizes a well-recognized intuition — that institutions must adapt at least as fast as the disruptions they govern — but converts it into an operationalizable threshold with policy implications that institutional lag theory does not itself yield. Institutional lag theory identifies that temporal asymmetry between technological change and institutional adaptation creates periods of maladjustment (Ogburn, 1922; North, 1990). The Compression–Absorption Framework specifies when that asymmetry becomes a structural risk, through what mechanism it generates lock-in, and what governance trigger — $C(t)$ persistently exceeding $A(t)$ — should activate preemptive response. Persistence is not left to discretionary interpretation. Jurisdictions must define *ex ante* temporal and magnitude parameters for activation—such as rolling multi-year averages, sector-specific compression tolerances, or cumulative $G(t)$ thresholds—to prevent political manipulation. The framework does not prescribe uniform duration; it requires that persistence criteria be specified transparently prior to deployment escalation. The formalization is not a restatement of lag theory; it is its operationalization as an anticipatory governance instrument.

Conventional regulatory models respond to visible harm. Financial regulation tightened after crisis; environmental standards followed demonstrable damage; data protection regimes expanded after misuse scandals. This reactive model assumes institutional correction can occur before structural reorganization becomes irreversible (North, 1990; Pierson, 2000) — an assumption that holds when technological diffusion unfolds gradually. Labor-compressing AI systems violate this assumption. When deployment scales rapidly across sectors, organizational hierarchies, capital allocation, and decision processes may reorganize faster than retraining systems, fiscal buffers, and regulatory frameworks can adjust. Increasing returns dynamics suggest that early architectural choices become progressively resistant to reversal (Arthur, 1989; David, 1985), narrowing the window for low-cost corrective intervention.

Institutional preemption follows directly from the framework’s logic: if $C(t)$ persistently exceeds $A(t)$ during deployment acceleration, cumulative lock-in pressure $P(T)$ rises — and the cost of later correction rises with it. The governance imperative is therefore to act during the preemption window: the period during which $G(t)$ is observable but $P(T)$ has not yet crossed entrenchment thresholds. The three structural commitments of preemption — early classification, threshold-triggered review, and compression monitoring — are designed to keep governance inside that window.

3.6 Retrospective Validation: Industrial Robots, 1990–2015

Before applying the framework to contemporary AI, a retrospective demonstration against a documented historical episode is instructive. Acemoglu and Restrepo (2020) provide the most rigorous available estimates of robot-driven labor displacement in the United States, covering the period 1990–2015. Their commuting-zone analysis quantifies the employment and wage effects of robot adoption across sectors. This body of evidence permits a method

demonstration of $C(t)$ and $A(t)$ construction under historical conditions — one where outcomes are known and the framework can be assessed against them.

Compression proxy $C(t)$: Acemoglu and Restrepo (2020) estimate that one additional robot per thousand workers reduced the employment-to-population ratio by approximately 0.34 percentage points and wages by 0.42 percent (published JPE version, Table 3 preferred specification). Between 1993 and 2007, US manufacturing absorbed approximately 400,000 industrial robots (net additions to US manufacturing robot stock, 1993–2007, consistent with IFR data cited in Acemoglu & Restrepo, 2020), concentrated in automotive, metals, and plastics sectors. Using task substitution rates derived from occupation-level exposure scores — the authors’ “exposure-to-robots” variable — a compression index can be constructed as the percentage of the affected workforce experiencing task substitution per year, weighted by wage displacement. During peak adoption (2000–2007), this index yields $C(t)$ values in the range of 1.2–2.0 percent of manufacturing workforce per year in high-exposure commuting zones.

Absorption proxy $A(t)$: For the same period, US institutional absorption was structurally limited. Trade Adjustment Assistance (TAA) served approximately 60,000–80,000 workers annually — a fraction of the displaced population. Cross-sector reemployment rates for displaced manufacturing workers remained below 60 percent within two years, with wage losses averaging 15–20 percent for those who did reemploy (Davis & von Wachter, 2011). Fiscal redeployment was primarily through Disability Insurance, Social Security, and unemployment insurance — income replacement rather than structural absorption. $A(t)$ estimates for US manufacturing circa 2000–2007 therefore range from 0.6 to 0.9 percent of affected workforce per year.

Framework application: During peak robot adoption, $C(t) \geq A(t)$ in high-exposure zones, yielding positive $G(t)$ values concentrated in the US Midwest and South. Acemoglu and Restrepo’s (2020) documented outcomes — persistent employment decline, wage scarring, and no reemployment convergence by 2015 — are consistent with cumulative lock-in pressure $P(T)$ having crossed structural entrenchment thresholds in affected commuting zones before corrective policy scaled. The retrospective application does not assert that the framework would have perfectly predicted outcomes; baseline $A(t)$ estimates are approximated from reemployment and program enrollment data not originally collected for this purpose. It demonstrates that the framework’s diagnostic logic, applied retroactively, correctly identifies the period and sectors where institutional absorption failed to keep pace with compression velocity — which is precisely the pattern the literature independently documents. This consistency is the relevant validation standard for a structural diagnostic framework rather than a predictive model.

4. Institutional Preemption: Governing Velocity Before Entrenchment

This approach differs from prevailing AI governance frameworks. Contemporary regulatory models—such as the OECD AI Principles and the EU AI Act—focus primarily on model-level risks, including safety, transparency, bias mitigation, and fundamental rights compliance. These frameworks classify systems according to risk categories based on potential harm at the point of use. By contrast, institutional preemption introduces a macro-institutional variable: deployment velocity relative to adaptive capacity. The concern is not solely whether a system is safe or fair in isolation, but whether cumulative adoption at scale generates structural compression before institutional recalibration becomes feasible.

Institutional preemption therefore involves three structural commitments.

First, early classification of high-autonomy systems capable of cross-sector decision-layer substitution. Classification is not prohibition. It is a risk-tiering mechanism that identifies systems whose aggregate deployment may influence labor share, entry pathways, or authority distribution at scale. Path dependency literature underscores the importance of early categorization before entrenchment effects accumulate (Pierson, 2000).

Second, structured review triggered by deployment scale or velocity thresholds. When projected compression in a sector exceeds defined benchmarks relative to absorptive indicators, phased rollout or conditional approval mechanisms may be warranted. This graduated response reflects precautionary governance logic under asymmetric risk conditions (Sunstein, 2005; Taleb, 2012), without imposing blanket restrictions.

Third, integration of macro-level compression monitoring into AI oversight regimes. Task-based labor research demonstrates that substitution effects often manifest within occupations before appearing in aggregate unemployment statistics (Autor, 2013; Acemoglu & Restrepo, 2019). Without monitoring task compression ratios, entry-level absorption rates, and wage dispersion trends, policymakers cannot assess whether deployment velocity is outpacing adaptation.

Institutional preemption thus reframes governance from reactive stabilization to anticipatory equilibrium management. Its aim is preservation of corrective optionality: acting prior to visible crisis, before $P(T)$ crosses entrenchment thresholds, preserves governance flexibility at lower cost than reconstruction after lock-in.

Importantly, preemption does not freeze institutional evolution. Economic systems historically adapt to technological transformation through new forms of participation and organizational design (Goldin & Katz, 2008; Brynjolfsson et al., 2021). What must be preserved is not a particular employment structure, but the capacity of institutions to recalibrate before compression hardens into constraint.

5. Human Agency and Authority Preservation

Labor compression is not solely an economic phenomenon. When labor-compressing AI systems automate decision-layer functions, three analytically distinct risks emerge: labor compression risk (the macro-economic displacement dynamic formalized in Section 3), authority concentration risk (structural redistribution of decision power away from accountable officials), and accountability deficit risk (erosion of the procedural mechanisms through which error is detected and corrected). These risks are causally connected but require separate analytical treatment to generate precise governance responses.

5.1 Labor Compression Risk

The labor compression risk is formalized as $C(t)$ in Section 3. At the sectoral level, task substitution within decision-layer occupations — analytical roles, managerial coordination, regulatory adjudication — reduces labor demand, narrows entry pathways, and concentrates returns to system operators. This economic dimension is addressed by the detection, modulation, and absorption architecture in Section 6. The connection to authority is structural: as labor compression reduces the number of human decision-makers within institutions, it also reduces the distributed pool from which accountability can be derived. Economic displacement and authority redistribution are therefore co-produced by the same deployment dynamic, but they require different policy instruments.

5.2 Authority Concentration Risk

Modern states depend on identifiable chains of authority and accountable decision-making structures (North, 1990; Fukuyama, 2014). Authority migration occurs when decision-layer processes in domains such as credit allocation, public benefit eligibility, regulatory enforcement, hiring, or healthcare triage become predominantly automated. Even where human review formally exists, override capacity may be limited by institutional incentives, performance metrics, or opacity of model outputs (Eubanks, 2018; Pasquale, 2015; Zuboff, 2019). Authority shifts from distributed human agents to centralized system operators — and in cases where those operators are transnational data platform companies, to actors whose extraction logic Zuboff (2019) characterizes as structurally incompatible with distributed democratic governance. In globally integrated digital markets, those operators may be headquartered beyond national jurisdiction — a sovereignty constraint identified by Rodrik (2011) in the context of economic integration and extended here to computational authority.

This risk is analytically distinct from both labor displacement and bias concerns. A system can be demonstrably unbiased in its output distribution and fully compliant with EU AI Act risk-tier requirements, yet still concentrate decision authority in ways that erode distributed governance capacity. A necessary qualification: the reviewer's question of whether computational mediation is meaningfully different from bureaucratic mediation is well-posed. Traditional bureaucracies also concentrate decision authority, limit discretion through procedural rule, and exhibit opacity in complex case-handling. The distinction this paper asserts is empirical rather than categorical. Algorithmic systems differ from bureaucratic mediation in three specific respects that are relevant to accountability structure: they operate at speed and scale that make sequential human review structurally impractical rather than merely costly; their decision logic is often proprietary and not subject to the freedom-of-information and administrative law discovery mechanisms that apply to human bureaucratic

records; and their update cycles can alter decision criteria without the procedural transparency of regulatory amendment. These are differences of degree that cross a governance-relevant threshold in certain deployment configurations, not a categorical distinction between computational and human authority. Whether they cross that threshold in specific cases is an empirical question, and the framework makes no prior assumption that all automation produces authority concentration. The concern is structural: authority concentration in computational systems reduces the redundancy and contestability that distributed human decision-making provides. Economic participation has long been associated with political legitimacy (Lipset, 1959); as decision-layer authority migrates from distributed officials to centralized algorithmic architectures, the institutional foundations of democratic legitimacy may weaken even without overt political disruption.

5.3 Accountability Deficit Risk

The accountability deficit risk concerns the erosion of procedural mechanisms through which governance error is detected, attributed, and corrected — a concern rooted in administrative law and institutional theory rather than labor economics or democratic theory. The operative mechanism is institutional hollowing: responsibility remains legally assigned to human actors yet meaningful discretion is constrained by automated systems whose internal logic is not transparently contestable (O’Neil, 2016). Over time this erodes procedural accountability and narrows avenues for corrective intervention — outcomes measurable through audit failure rates, appeal overturn rates, and the proportion of consequential decisions subject to documented human review.

This argument differs from general calls for “human-in-the-loop” oversight. Symbolic human presence without meaningful discretion does not preserve institutional accountability. The safeguard must ensure substantive review capacity and liability attachment: identifiable officials in authority-sensitive domains — legal adjudication, immigration determination, welfare eligibility, financial access, regulatory sanction — must retain documented override capacity and traceable reasoning obligations. Auditability, version control, and appeal pathways are structural requirements for institutional durability, not optional ethical enhancements (North, 1990). The accountability deficit risk interacts with the authority concentration risk: when decision authority concentrates in opaque systems, the procedural mechanisms designed to detect and correct error in distributed human decision-making become inapplicable, and no equivalent mechanism is automatically substituted.

5.4 Analytical Integration

The three risks interact but require distinct governance responses. Labor compression risk is addressed by the velocity modulation and absorption scaling mechanisms in Section 6. Authority concentration risk requires sectoral classification and human-decision-authority mandates in legitimacy-critical domains — the Human Agency Safeguard Protocol in Appendix C. Accountability deficit risk requires procedural design: override mechanisms, auditability requirements, and appeal infrastructure that ensure consequential decisions remain contestable.

The proportionality principle governs all three: autonomy constraints scale with domain sensitivity, and industrial optimization, logistics automation, and scientific modeling do not implicate legitimacy in the same way as rights-determining functions. Human agency in this framework is institutional rather than metaphysical — it refers to the capacity of governance systems to maintain accountable decision structures and recalibrate when errors accumulate, sustaining corrective optionality in the face of accelerating deployment.

6. Policy Architecture: Detection, Modulation, and Absorption

Operationalizing institutional preemption requires layered governance architecture consistent with risk management theory and institutional adaptation scholarship (North, 1990; Sunstein, 2005).

6.1 Detection Layer

The detection layer establishes systematic monitoring of compression dynamics. Large-scale deployments of high-autonomy systems should trigger disclosure of anticipated task substitution effects and sectoral impact. Task-based labor research demonstrates that automation effects often unfold within occupations before manifesting as aggregate displacement (Autor, 2013; Acemoglu & Restrepo, 2018). Aggregated data enable periodic Compression–Absorption assessments at the national level, allowing policymakers to observe whether velocity differentials are emerging. Transparency in such monitoring enhances legitimacy and facilitates democratic oversight, reinforcing institutional accountability structures (Lipset, 1959; Fukuyama, 2014). Developing nations and smaller ASEAN member states need not build detection infrastructure from scratch. Several open-source and standards-based tools can reduce startup costs. The OECD.AI Policy Observatory provides a freely accessible framework of indicators aligned with the OECD AI Principles that can be adapted for national compression monitoring. ISO/IEC 42001:2023 (AI Management Systems) establishes baseline audit and documentation requirements compatible with the detection layer’s disclosure obligations. The ILO’s occupational task database (O*NET-aligned ISCO mapping) provides a publicly available foundation for constructing task substitution indices. For fiscal absorption indicators, IMF Fiscal Monitor templates provide standardized formats that many member states already report. Integrating these existing frameworks into the detection layer reduces implementation barriers while enabling cross-national comparability of $G(t)$ estimates.

In labor markets with substantial informal employment, official displacement and wage statistics may understate compression dynamics. Informal-sector absorption can temporarily buffer displaced workers without reflecting compression in formal employment indicators, creating false negatives in $C(t)$ estimation. Detection frameworks in such contexts should therefore incorporate proxy indicators—such as underemployment rates, earnings volatility among self-employed workers, sectoral income dispersion, and shifts in informal-to-formal transition patterns—to reduce measurement bias. The objective is not perfect precision, but conservative detection sufficient to prevent velocity asymmetry from remaining statistically invisible.

6.2 Modulation Layer

The modulation layer introduces velocity controls without blanket prohibition. Threshold-triggered review processes allow policymakers to evaluate whether deployment pace aligns with absorptive capacity before increasing returns dynamics harden structural reorganization (Arthur, 1989; Pierson, 2000). Sectoral differentiation is critical in open economies. Export-intensive industries operating under global competitive pressure may tolerate faster compression if productivity gains offset domestic labor impact, consistent with globalization constraints on national policy autonomy (Rodrik, 2011). By contrast, sectors anchoring civic

legitimacy—public administration, welfare systems, legal adjudication—require stricter thresholds to preserve authority structures and accountability chains central to institutional durability (North, 1990; Fukuyama, 2014). The procedural design of threshold-triggered review should follow three principles. First, triggers are defined *ex ante* and published: a deploying firm or regulatory agency knows in advance that when the detection layer’s $C(t)$ estimate exceeds the sector-specific threshold, a mandatory review is initiated — reducing regulatory uncertainty and preventing strategic deployment timing. Second, the review is time-bounded (90–180 days is consistent with comparable precautionary instruments such as REACH chemical assessments and FSOC bank stress reviews), producing either conditional approval with absorption-scaling requirements, phased rollout consent, or deferral pending $A(t)$ evidence. Third, the burden of demonstrating absorption sufficiency rests with the deploying entity, not the regulator — a precautionary inversion consistent with Article 191 TFEU environmental logic applied to macro-institutional risk.

The escalation pathway implied by the modulation layer is operational rather than discretionary. Detection of $G(t) > 0$ initiates tiered monitoring, duration assessment, and cross-domain persistence evaluation prior to stability intervention. The structured decision sequence is summarized in Figure 1.

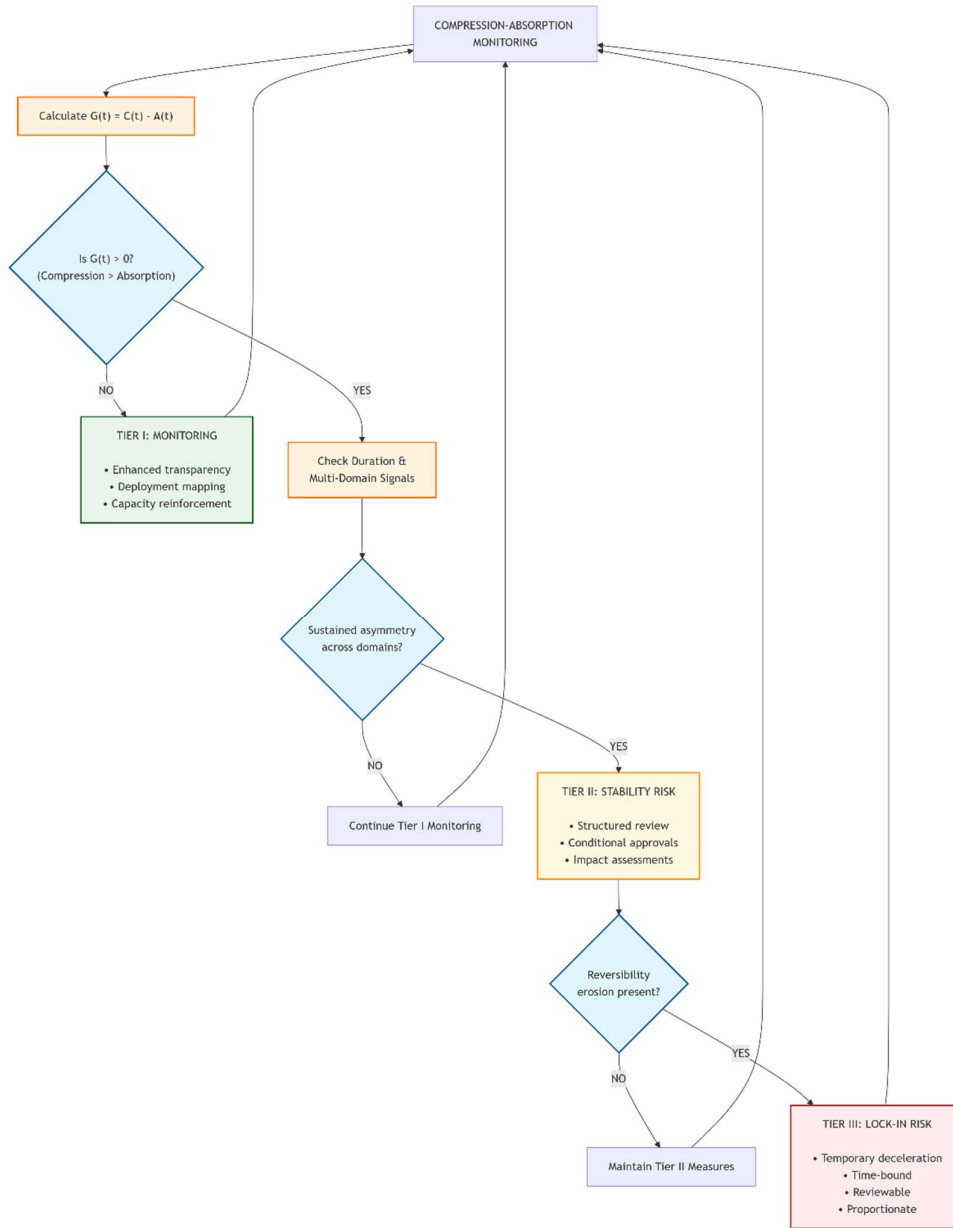


Figure 1. Tiered Compression–Absorption Escalation Framework.

Rule-based monitoring and escalation logic linking measured compression–absorption differentials $G(t)$ to graduated policy responses (Tier I–III) based on duration, cross-domain persistence, and reversibility risk.

Indicator	Tier I – Monitoring	Tier II – Stability Risk	Tier III – Lock-In Risk
G(t) Status	$G(t) \leq 0$ or transient exceedance	Sustained $G(t) > 0$ above defined sector threshold	Escalating and structurally elevated $G(t)$
Duration (Defined Ex Ante)	Short-term fluctuation within defined window	Multi-period exceedance across rolling window	Persistent structural exceedance beyond defined persistence threshold
Cross-Domain Scope	Single-sector or localized impact	Multi-sector diffusion	Cross-systemic asymmetry across core economic domains
Reversibility Assessment	High optionality	Moderate erosion of optionality	Declining or materially eroded correction window
Governance Response	Enhanced monitoring and transparency	Structured review and conditional modulation	Temporary deceleration or phased deployment subject to review
Activation Mechanism	Automatic reporting and disclosure	Threshold-triggered review defined ex ante	Escalation protocol defined ex ante; proportionate and reviewable

Table 6.1 Illustrative Trigger Tier Calibration Matrix: Jurisdictions define persistence, magnitude, and cross-domain parameters ex ante to reduce discretionary activation while preserving contextual flexibility.

6.3 Absorption Layer

The absorption layer strengthens adaptive capacity. Scaling retraining pipelines aligned with emerging industries, enhancing labor mobility, investing in human-centered service sectors, and modernizing fiscal buffers all increase $A(t)$. General-purpose technology literature suggests that productivity gains translate into broad-based prosperity only when accompanied by complementary institutional investment (Brynjolfsson et al., 2021). Importantly, absorption capacity is endogenous; innovation can expand it. The framework therefore does not assume pessimistic outcomes. Rather, it requires that adaptation scale alongside compression, consistent with dynamic adjustment theory (Acemoglu & Restrepo, 2018).

6.4 Political Economy of Institutional Preemption

The policy architecture outlined above is technically coherent. Its political viability is less certain. The analysis in Sections 6.4.1–6.4.3 is explicitly conjectural coalition modeling rather than empirically grounded political economy. It draws on structural features of comparable governance episodes — macroprudential regulation, environmental impact assessment, and gig economy legislation — to generate plausible resistance and support configurations. It does not present original comparative evidence, and the configurations it identifies are context-dependent: they may hold in middle-income open economies with organized civil service sectors but are not claimed to generalize to political systems with different institutional configurations. Readers should treat this section as a design framework for political feasibility analysis rather than an empirical prediction of coalition behavior. With that scope condition stated, the structural challenge is as follows. Institutional preemption imposes immediate, concentrated costs on powerful actors while generating diffuse, probabilistic benefits for future stability. This temporal and distributional asymmetry creates identifiable resistance configurations that velocity governance design must address.

6.4.1 Resistance Dynamics and Capture Risk

Multiple constituencies have structural incentives to oppose or dilute velocity governance:

Deployers of labor-compressing AI systems—large firms, platforms, and major employers—prioritize rapid deployment to capture market share, reduce labor costs, and establish network effects before competitors. Resistance manifests through competitiveness narratives ("regulation will make us uncompetitive globally"), lobbying against disclosure requirements and velocity thresholds, and jurisdictional arbitrage toward friendlier regulatory environments. These actors typically possess substantial political influence through campaign contributions, revolving-door relationships, and technical expertise asymmetries that can overwhelm under-resourced regulators.

Export-exposed manufacturers and tradable service providers fear that velocity controls will impose cost disadvantages relative to competitors in jurisdictions without such governance. Their resistance takes the form of relocation threats, arguments that "Malaysia cannot regulate ahead of the US or China," and claims that any modulation will trigger capital flight. This constituency wields particular influence in open economies dependent on foreign direct investment and global value chain integration.

Finance and investment sectors favor labor-light growth models that promise scalable margins and "AI premium" valuations. Resistance focuses on opposing productivity-linked

contributions to absorption systems, preferring voluntary codes and self-regulation. Financial actors can threaten to downgrade sovereign creditworthiness or redirect investment flows, creating fiscal pressure on governments.

State capacity skeptics within government may resist institutional preemption to avoid responsibility for measurement, enforcement, and potential blame if economic growth slows. Bureaucratic resistance manifests as procedural delay ("more studies needed"), scope narrowing (limiting policy to model safety rather than velocity governance), and systematic underfunding of monitoring infrastructure. Career incentives often favor risk-averse inaction over potentially controversial anticipatory intervention.

Politically connected intermediaries may seek to extract rents from procurement processes, vendor selection, and licensing mechanisms. Their interest lies in converting velocity governance into patronage opportunities through vague licensing requirements rather than substantive compression monitoring. This dynamic is particularly acute in contexts where regulatory capture and rent-seeking are endemic.

Professional elites in middle-management and credentialed gatekeeping roles exhibit split incentives. Some seek protection from automation, while others resist auditability requirements that expose discretionary decision-making or reduce information rents. This constituency's political alignment depends on whether they perceive velocity governance as protecting their occupational status or threatening their autonomy.

The core political economy challenge is straightforward: institutional preemption imposes early, certain costs (disclosure systems, audit capacity, phased rollout requirements) for probabilistic future benefits (avoiding structural lock-in). Concentrated costs mobilize opposition; diffuse benefits struggle to generate equivalent political support. This is the classic logic of regulatory capture (Stigler, 1971).

6.4.2 Supportive Constituencies and Coalition Architecture

Despite formidable resistance, institutional preemption can mobilize countervailing support from constituencies whose interests align with velocity governance:

Labor-exposed groups—including unions where organized, civil service associations, professional bodies with licensing traditions (medicine, law, accounting, engineering), and education sector organizations—face direct compression risk. These actors possess organizational capacity, electoral influence in some jurisdictions, and normative legitimacy in debates over economic inclusion. Their political effectiveness depends on density of organization and institutional embeddedness.

Domestic small and medium enterprises (SMEs) are often disadvantaged by unfettered automation because large firms can adopt high-autonomy systems faster and use resulting cost advantages to undercut smaller competitors. SMEs may support institutional preemption if it prevents "winner-take-most" deployment races that consolidate market power. This constituency is politically significant in economies where SMEs provide substantial employment and represent geographically distributed interests.

Social protection institutions—including agencies like Malaysia's SOCSO, Human Resource Development Corporation (HRD Corp), and Ministry of Human Resources

(MOHR)—directly experience fiscal and administrative burdens from labor market disruption. These institutions have bureaucratic incentives to support early intervention that prevents crisis-driven surges in demand for services. Their technical expertise in labor market dynamics can provide credible policy justification.

Public-interest regulators—central banks concerned with financial stability, competition authorities monitoring market concentration, data protection agencies, and sector-specific regulators (health, finance, employment)—may recognize that unchecked velocity creates systemic risks within their mandates. Cross-cutting regulatory concerns can overcome siloed resistance if coordination mechanisms exist.

"Human-in-the-loop" AI vendors selling augmentation tools rather than full automation systems have commercial interests aligned with governance that preserves human authority in decision-layer functions. This constituency can provide technical counterweight to full-automation advocates and demonstrate commercial viability of alternative deployment models. The coalition architecture outlined above is analytically plausible but empirically underspecified at this stage. Each constituency's alignment is conditional rather than structural: labor organizations may prioritize productivity-sharing over deployment constraints; SMEs may prefer regulatory exemptions over sector-wide velocity limits; fiscal institutions may face austerity pressures that override preemption logic. Counterexamples exist — in the United Kingdom, trade unions have supported automation where productivity agreements guaranteed wage sharing; in Singapore, government-led automation has proceeded with limited SME resistance due to tight labor markets. The political feasibility model presented here is illustrative of a design logic, not a prediction of coalition behavior. Whether and under what conditions these alignments materialize requires country-specific political economy analysis beyond the scope of this paper.

6.4.3 Designing Politically Viable Velocity Governance

Political feasibility requires coalition design that neutralizes opposition while mobilizing support:

Anchor in authority-sensitive domains. Framing velocity governance around legitimacy-critical functions—legal adjudication, welfare eligibility determination, credit access, healthcare triage, public administration—shifts the debate from job protection to institutional integrity. This reframing appeals to rule-of-law constituencies, accountability advocates, and sovereignty-concerned actors who might otherwise dismiss labor-focused arguments.

Provide export sectors an "acceleration lane." Explicitly permitting faster deployment in tradable goods and services—with clear reporting requirements and predictable contributions to absorption systems—addresses competitiveness concerns while maintaining governance oversight. This sectoral differentiation, outlined in Section 6.2, transforms potential opponents into conditional supporters.

Prevent predatory scale advantages through disclosure. Transparency requirements and phased rollout processes reduce information asymmetries that allow large firms to capture first-mover advantages before smaller competitors can respond. Framed as fair competition rather than technological restraint, this appeals to SMEs and competition authorities.

Establish a fiscal bargain tied to productivity gains. Rather than ad-hoc taxation, productivity-linked contributions to absorption systems (retraining, mobility support, transition assistance) create predictable cost structures. Framing contributions as complementary investment in human capital—necessary for productivity gains to translate into broad-based prosperity—aligns with general-purpose technology literature (Brynjolfsson et al., 2021) and positions velocity governance as growth-enabling rather than growth-constraining.

Build on existing anticipatory governance precedents. Financial capital buffer requirements, environmental impact assessments, and public health systems all institutionalized risk mitigation prior to visible crisis (Hacker, 2004; Sunstein, 2005). Analogizing velocity governance to these accepted frameworks reduces perceived novelty and leverages established legitimacy of precautionary approaches.

Political feasibility therefore depends on coalition strategies that transform institutional preemption from "job protection" into "fair competition, institutional legitimacy, and sustainable growth." The framework becomes politically sustainable when framed as preserving both competitiveness and systemic stability rather than choosing between them.

6.5 Integration and Feasibility

Governments often act after crisis rather than before, reflecting well-documented patterns of reactive institutional reform (Pierson, 2000). The political economy obstacles outlined above explain this tendency. Yet anticipatory governance is not unprecedented.

In growth-constrained developing states, competitive pressures may incentivize acceleration even when $G(t)$ rises. This dynamic increases the importance of institutionalized, rule-based trigger mechanisms rather than reliance on discretionary political judgment. Where growth competition is intense, velocity governance must be embedded in procedural safeguards rather than contingent on executive restraint.

Where competitive pressures override threshold activation, the framework predicts accelerated accumulation of $P(T)$ rather than spontaneous institutional rebalancing.

Financial capital buffer requirements were institutionalized to reduce systemic fragility before recurrence of collapse (Hacker, 2004; Sunstein, 2005). Environmental impact assessments and public health vaccination systems similarly embed risk mitigation prior to visible catastrophe. Institutional preemption extends this logic to labor-compressing AI under asymmetric risk conditions (Taleb, 2012).

By integrating detection, modulation, and absorption layers—while addressing political economy constraints through strategic coalition-building—states can manage velocity without suppressing innovation. The objective is not technological restraint, but alignment between acceleration and institutional equilibrium, preserving both competitiveness and systemic stability (North, 1990; Rodrik, 2011). Political viability emerges not from technocratic imposition, but from coalitions that recognize shared interests in preventing velocity asymmetry from hardening into structural constraint.

7. Normative Clarifications and Scope Conditions

7.1 Scope Conditions

Four scope conditions bound the framework’s claims. First, the argument is structural, not predictive: the Compression–Absorption Framework does not assert that catastrophic unemployment is imminent or inevitable. It identifies a conditional stability requirement—when $C(t)$ persistently exceeds $A(t)$, structural lock-in risk accumulates—consistent with institutional lag theory (North, 1990; Pierson, 2000) and asymmetric risk reasoning (Sunstein, 2005; Taleb, 2012). Second, institutional preemption is not technological restraint: the doctrine targets deployment velocity, not innovation direction. It does not shield inefficient sectors from competition; it recognizes that macro-level compression can generate cross-sector feedback effects that outpace institutional recalibration, particularly in open economies (Rodrik, 2011). Third, human agency in this framework is institutional, not metaphysical: the concern is that when decision-layer authority migrates into opaque computational systems, identifiable accountability chains weaken and corrective optionality narrows—not that automation is inherently dehumanizing. Fourth, the framework does not claim to model the precise functional relationship between cumulative compression pressure and lock-in probability; empirical calibration across sectors and jurisdictions remains an open research question addressed in §6.2.

7.2 Empirical Falsification: What Would Disprove the Framework

The Compression–Absorption Framework advances structural hypotheses that must remain empirically testable. This section specifies observable conditions that, if demonstrated over a multi-year period following widespread AI deployment, would falsify or substantially weaken the framework's core predictions. These conditions cannot be evaluated from current data alone; they require longitudinal observation across jurisdictions after labor-compressing AI systems have been deployed at scale. Clear falsification criteria distinguish diagnostic frameworks from unfalsifiable ideology and enable productive empirical research.

Falsifier Set A: Velocity Does Not Create Governance Gaps

The framework predicts that when compression accelerates beyond absorptive capacity, observable labor market strain emerges through specific indicators.

Falsification Condition A.1: Across multiple sectors and jurisdictions, rapid adoption of high-autonomy systems occurs without:

- Persistent rise in youth unemployment or underemployment rates relative to baseline
- Wage polarization between high-skill and mid-tier analytical occupations
- Declining entry pathways into decision-layer professional roles
- Pressure on labor share of national income
- And simultaneously with stable or improving mobility indicators (retraining completion leading to reemployment within 6–12 months at comparable wage levels)

If fast deployment velocity consistently correlates with broad-based wage stability, sustained labor force participation, and smooth occupational transitions over a multi-year observation window following widespread AI deployment, the compression-absorption gap mechanism

would be falsified or require substantial revision. These conditions would need to persist across multiple sectors and jurisdictions, not merely reflect short-run labor market tightness. The framework would be weakened if velocity proves irrelevant to labor market outcomes.

Falsifier Set B: Path Dependency and Lock-In Do Not Materialize

The framework predicts that increasing returns dynamics make structural reorganization around high-autonomy systems progressively difficult to reverse.

Falsification Condition B.1: Countries or sectors that initially allow high-autonomy deployment at scale can later reverse or substantially reshape automation architectures at low economic and political cost, including:

- Re-expanding human decision-layer functions in previously automated domains
- Restoring substantive (not merely formal) human override authority
- Accomplishing such reversal without major productivity losses, capital flight, or political crisis

If reversibility proves cheap, common, and politically uncontested—contradicting increasing returns predictions—the "irreversibility window" concept would be invalidated. Observing frequent, low-cost institutional recalibration after high-velocity deployment would undermine the preemption rationale.

Falsifier Set C: Authority Migration Does Not Concentrate Power or Erode Accountability

The framework posits that decision-layer automation can centralize authority in ways that weaken institutional accountability and corrective capacity.

Falsification Condition C.1: Even with extensive AI deployment in authority-sensitive domains (legal adjudication, credit allocation, welfare eligibility, regulatory enforcement), we observe:

- Strong transparency mechanisms and consistently auditable decision trails
- High rates of successful appeals and meaningful review processes
- Clear liability attachment to identifiable human decision-makers
- No measurable decline in public trust or perceived institutional legitimacy
- Distributed rather than concentrated influence over algorithmic governance

If AI integration systematically improves accountability, enhances transparency, and distributes authority more broadly than pre-automation systems, the sovereignty and authority-migration concerns would be weakened. The framework assumes computational centralization can create accountability deficits; evidence of the opposite would require theoretical revision.

Falsifier Set D: Absorption Scales Endogenously Faster Than Compression

The framework permits absorptive capacity $A(t)$ to expand but expresses concern that it often lags compression velocity $C(t)$.

Falsification Condition D.1: Across repeated episodes and diverse jurisdictions, governments demonstrate capacity to scale absorption mechanisms faster than compression advances ($A(t) \geq C(t)$ persistently) without unsustainable fiscal stress, including:

- Retraining systems that successfully redeploy displaced workers at scale
- Labor mobility mechanisms that facilitate rapid cross-sector transitions
- Wage insurance and income stabilization that prevent prolonged displacement
- Fiscal capacity to sustain such programs without undermining macroeconomic stability

If the median case proves to be "institutions keep pace automatically," the institutional preemption doctrine becomes optional risk management rather than structural necessity. The framework would require recalibration if absorption routinely matches or exceeds compression velocity endogenously.

Methodological Implications

These falsification conditions are deliberately specific and testable using labor market data, institutional response metrics, and comparative case studies. The framework does not claim:

- That all automation generates governance gaps (only when velocity substantially exceeds capacity)
- That lock-in is inevitable (only that it becomes progressively costly to reverse)
- That authority always centralizes (only that it may do so absent safeguards)
- That absorption never keeps pace (only that it frequently lags)

Rather, the Compression–Absorption Framework predicts conditional instability when velocity exceeds capacity by substantial margins over sustained periods. Observing the opposite—rapid adoption consistently accompanied by broad-based stability, easy reversibility, distributed accountability, and endogenous absorption scaling—would necessitate framework revision or abandonment.

This testability distinguishes the framework from non-falsifiable advocacy positions. Empirical research demonstrating any of the above falsification conditions — observed over a sustained multi-year period following widespread AI deployment, not from current snapshot data — would constitute legitimate grounds for rejecting or substantially modifying the framework's predictions.

A final methodological complication deserves direct acknowledgment: the preemption paradox. If institutional preemption functions as intended — if detection, modulation, and absorption mechanisms prevent $G(t)$ from accumulating into lock-in — the post-hoc empirical record will show labor market stability rather than crisis. A skeptical observer may then conclude that preemption was unnecessary because no crisis materialized. This is the standard epistemological problem of preventive policy evaluation (Sunstein, 2005): successful prevention is observationally equivalent to unnecessary intervention. The framework cannot escape this constraint.

Its response is asymmetric risk logic: the cost of institutional preparation that proves unnecessary is modest; the cost of delayed preparation once lock-in has advanced is substantially higher. If the framework's falsifiers are met — if widespread AI deployment does not generate persistent $C(t) > A(t)$ — the correct interpretation is that the stability condition was satisfied, not that the framework was irrelevant. The framework's empirical leverage operates prospectively as deployment scales, not retrospectively from current data — which is precisely what the falsification conditions above are designed to enable.

8. Regional Application: Open Economies and ASEAN Context

8.1 Structural Position of Open Mid-Sized Economies

The empirical literature on AI-driven labor displacement has developed almost entirely within advanced economy contexts — the United States, European Union, and OECD members — where institutional absorptive capacity is deep, labor markets are highly formalized, and displacement signals surface through measurable unemployment and wage data. This paper engages a structurally different context. The governance challenge in developing and middle-income economies, particularly across ASEAN, is not identical to the challenge that literature documents. Three structural differences are analytically consequential and shape the application of the Compression–Absorption Framework in this section.

First, the exposure profile is inverted. In advanced economies, the highest AI-exposed roles are high-income, high-education decision-layer occupations — workers with individually strong adaptive capacity. In the East Asia and Pacific region, only approximately 10 percent of jobs involve tasks complementary to AI, compared to 30 percent in advanced economies, and the roles most vulnerable to substitution are those integrated into global outsourcing architectures — business process outsourcing, shared services, back-office data processing, and call centre functions (Arias, Fukuzawa, Le, & Mattoo, 2025). These are precisely the formal sector entry pathways through which middle-income economies have channelled youth into the labor market. Compression of these roles is not professional hollowing of high-income workers; it is destruction of the upward mobility ladder.

Second, the absorptive mechanism differs fundamentally. In advanced economies, displacement generates observable unemployment and wage compression signals that policy systems can detect and respond to. In ASEAN labor markets, displacement from formal sector roles is absorbed into the informal sector rather than measured unemployment. Arias et al. (2025), in the most rigorous causal study of ASEAN robot adoption to date covering 2018–2022, found that robot adoption displaced approximately 1.4 million low-skilled formal workers across five ASEAN countries — but these workers did not appear in unemployment statistics; they transitioned to informal employment. This informal-sector buffer makes $G(t) > 0$ structurally invisible in official data. Headline unemployment stability, such as Malaysia’s ≈ 3.5 percent documented in Section 9, is therefore not a falsification of compression pressure; it is a measurement artifact of the absorption mechanism.

Third, the locus of the compressing decision is extraterritorial. In advanced economies, firms deploying labor-compressing AI are primarily domestic or subject to domestic regulatory jurisdiction. In ASEAN, a substantial share of compression pressure originates with multinational firms making deployment decisions in jurisdictions entirely outside ASEAN regulatory reach — automating the functions of their regional outsourcing centres, shared service hubs, and supply chain management operations from headquarters in North America or Europe. The compressing entity and the compressed workforce are in different jurisdictions. This is a distinct governance problem that velocity modulation frameworks designed for domestic deployment cannot address unilaterally — and it is one of the structural arguments for regional coordination developed in Section 8.3.

These three structural differences — inverted exposure profile, informal-sector buffering, and extraterritorial compression locus — define why developing-economy velocity governance requires frameworks calibrated to these conditions rather than transpositions of advanced-economy regulatory models. The sections that follow develop the Two-Speed Governance Model and ASEAN coordination architecture with these structural features as the design constraints.

With these structural differences as the design context, the governance challenge for open mid-sized economies can be stated more precisely. Open mid-sized economies confront acceleration asymmetry under distinct constraints. Unlike technological superpowers, they do not define global AI standards. Yet they are highly exposed to global deployment pressures and capital mobility. Malaysia and ASEAN member states exemplify this structural position: export-oriented growth models, integration into global supply chains, and reliance on foreign technology platforms amplify both competitiveness incentives and exposure to rapid labor compression.

In such contexts, institutional preemption must reconcile two imperatives: preserving competitiveness and safeguarding systemic equilibrium. Uniform deceleration would risk marginalization in global markets. Conversely, unmodulated acceleration may intensify domestic compression beyond absorptive capacity.

8.2 The Two-Speed Governance Model

The Two-Speed Governance Model provides a pragmatic reconciliation. Export-critical sectors—such as semiconductor manufacturing, logistics optimization, and industrial automation—may tolerate accelerated deployment where growth offsets compression. Domestic legitimacy-critical sectors—public administration, welfare eligibility systems, credit access determination, healthcare allocation, and regulatory enforcement—require stricter governance thresholds to preserve authority structures and labor absorption functions.

This sectoral differentiation recognizes that open economies face asymmetric constraints. Export sectors compete in global markets where unilateral velocity modulation could disadvantage domestic firms relative to foreign competitors operating without such governance. By contrast, legitimacy-critical sectors anchor civic trust and institutional coherence, where authority migration or compression-induced instability threatens political sustainability regardless of export competitiveness (Lipset, 1959; Fukuyama, 2014). The two-speed model thus preserves competitive positioning while safeguarding institutional foundations.

The classification criterion for lane assignment follows directly from Appendix B’s authority-sensitivity taxonomy (B4). Systems whose primary deployment domain is an authority-sensitive function—legal adjudication, public benefits eligibility, financial access, healthcare triage, regulatory enforcement—belong in the stability preservation lane regardless of sector label. Systems deployed primarily in tradable goods production, logistics, or export-oriented services with no material footprint in authority-sensitive functions qualify for the competitive acceleration lane. Novel sectors not cleanly assignable to either category default to the stability preservation lane pending empirical assessment of their authority footprint. This criterion gives regulators a replicable classification logic rather than requiring case-by-case political judgment.

8.3 The Imperative and Challenge of Regional Coordination

Regional coordination becomes particularly important in open economies. If individual states impose velocity modulation while neighboring jurisdictions do not, regulatory arbitrage may undermine preemptive efforts. Firms facing stricter governance in one ASEAN member state could relocate operations, shift investment, or restructure supply chains to jurisdictions with lighter regulatory touch. This dynamic creates pressure toward regulatory convergence at the lowest common denominator—precisely the race-to-the-bottom outcome that undermines collective institutional stability.

ASEAN-level dialogue on high-autonomy classification standards, human accountability requirements, and compression monitoring metrics could mitigate destabilizing competitive races. Coordinated frameworks would establish regional baselines, reduce arbitrage incentives, and enable states to implement velocity governance without unilateral competitiveness penalties. Such coordination need not eliminate national sovereignty; rather, it reinforces collective stability in open markets.

8.3.1 Why ASEAN Coordination Has Proven Difficult

Despite clear benefits, ASEAN coordination on labor-technology governance has been limited. Understanding these obstacles enables more realistic institutional design.

The non-interference norm and sovereignty sensitivity. ASEAN's foundational principles emphasize non-interference in domestic affairs and respect for national sovereignty. Binding coordination on velocity governance could be perceived as external constraint on policy autonomy, triggering political resistance even where economic logic favors cooperation. Member states guard domestic labor market policy as core sovereign prerogative, making supranational governance frameworks politically contentious (Acharya, 2009).

Developmental divergence and heterogeneous interests. ASEAN spans dramatic variance in economic development, labor market structure, and institutional capacity. Singapore operates as advanced financial hub with sophisticated retraining infrastructure and high absorptive capacity. Cambodia, Laos, Myanmar, and Vietnam (CLMV countries) face fundamentally different labor market dynamics, with large agricultural and informal sectors, limited fiscal buffers, and nascent automation adoption. Middle-income states like Malaysia, Thailand, and Philippines occupy intermediate positions. This heterogeneity makes one-size-fits-all coordination impractical; what Singapore can absorb may overwhelm Indonesian or Philippine institutional capacity (Rodrik, 2011).

The extraterritorial compression problem.

A compression risk distinctive to ASEAN's economic architecture is that a substantial share of compression pressure originates with multinational firms making deployment decisions outside ASEAN regulatory jurisdiction entirely. Malaysia's business process outsourcing and shared services sector — which employs several hundred thousand workers in accounting, legal processing, customer management, and back-office data operations — exists because multinational firms chose to locate these functions in the region. When those same firms deploy AI to automate these functions, the deployment decision is made at headquarters in jurisdictions that ASEAN velocity governance cannot reach. The compressed workers are in ASEAN; the compressing decision is in North America or Europe. This asymmetry means

that ASEAN-level governance can modulate domestic deployment velocity but cannot directly constrain the primary driver of outsourced-sector compression. Regional coordination mechanisms therefore have limited direct leverage on this risk; the more effective response is absorption-side investment — accelerating redeployment capacity before compression materializes — rather than attempting to govern deployment decisions made extraterritorially.

Race-to-the-top is hard; race-to-the-bottom is easy. Economic theory predicts that regulatory competition under capital mobility tends toward lax standards rather than stringent coordination (Vogel, 1995). Each ASEAN member state fears being the "slow adopter" that loses foreign direct investment to neighbors offering faster, cheaper deployment environments. First-mover disadvantages in imposing velocity governance create collective action problems: all states might benefit from coordinated standards, yet each has individual incentive to defect by offering lighter regulation to attract investment. Overcoming this requires credible commitment mechanisms that ASEAN's consensus-based architecture struggles to provide.

Administrative capacity variance and enforcement asymmetry. Even if member states agree on common velocity governance frameworks, implementation would require measurement systems, audit capacity, and enforcement mechanisms that vary dramatically across the region. Singapore's well-resourced regulators can monitor compression indicators and enforce compliance; less-developed members may lack statistical infrastructure, technical expertise, or budgetary resources for comparable oversight. Coordination without enforcement becomes symbolic rather than substantive, undermining credibility and inviting free-riding.

Domestic political cycles and short-term incentives. Institutional preemption generates benefits that materialize over medium to long time horizons, while political costs appear immediately through resistance from deploying firms and competitiveness concerns. Politicians facing near-term electoral pressures or coalition fragility have weak incentives to accept binding regional commitments that impose current costs for future stability. This temporal mismatch between political accountability cycles and institutional preemption payoffs discourages leaders from championing regional coordination (Pierson, 2000).

8.3.2 WTO Compatibility: Scope Note

Velocity governance mechanisms—particularly threshold-triggered deployment review and sectoral rollout conditions—raise potential WTO compatibility questions, specifically regarding whether such instruments constitute non-tariff barriers under GATT Article III, TBT Agreement disciplines, or GATS market-access commitments. Full analysis of WTO compatibility is beyond the scope of this paper, which addresses macro-institutional governance design rather than international trade law. Provisionally, velocity governance mechanisms are most defensible under the GATS public morals and public order exceptions (Article XIV) and GATT Article XX(a) analogies. GATS Article XIV permits measures "necessary to protect public morals or public order" and "to secure compliance with laws or regulations"—categories into which velocity governance in legitimacy-critical domains may plausibly fall. Defensibility is strengthened when measures are applied on a non-discriminatory basis (applicable equally to domestic and foreign deployers, preserving national treatment and MFN obligations), are no more trade-restrictive than necessary, and include transparent procedural design as specified in Section 6.2. Future legal research should

examine whether WTO dispute settlement jurisprudence would accommodate anticipatory macro-institutional governance, and whether ASEAN's own trade architecture provides jurisdictional space outside WTO disciplines.

8.4 Toward Modular and Pragmatic ASEAN Cooperation

These obstacles do not render regional coordination impossible, but they necessitate modular, graduated approaches that respect sovereignty constraints while building toward deeper integration.

Phase I: Common definitions and reporting templates. ASEAN coordination should begin with low-sovereignty-cost measures: developing shared definitions of high-autonomy systems, standardized compression indicators (aligned with Section 6.1 detection frameworks), and voluntary reporting templates. This phase establishes conceptual common ground and measurement infrastructure without binding commitments. Precedents exist in ASEAN's approach to cybersecurity frameworks and data governance dialogues, where soft coordination precedes harder integration (Caballero-Anthony, 2020).

Phase II: Mutual recognition of audits and assessments. Once measurement systems mature, member states could develop mutual recognition agreements for high-autonomy system audits conducted in specific domains. A firm deploying AI-driven credit allocation systems across multiple ASEAN markets could undergo standardized audit in one jurisdiction, with results recognized by others. This reduces compliance costs, prevents duplicative processes, and creates network effects favoring participation without requiring supranational enforcement. Mutual recognition preserves sovereignty—each state decides whether to recognize others' audits—while enabling practical coordination.

Phase III: Coordinated thresholds in legitimacy-critical domains. Only after trust and institutional capacity develop through Phases I and II should ASEAN consider region-wide thresholds for velocity governance in authority-sensitive sectors. Initial coordination might focus on narrow domains with high sovereignty salience—such as public administration, legal adjudication, or welfare systems—where competitiveness concerns are minimal and institutional legitimacy concerns are maximal. Export sectors could remain under national jurisdiction, preserving flexibility while establishing baseline protection for civic functions.

Sectoral pilot programs. Rather than comprehensive frameworks, ASEAN could pilot velocity governance coordination in specific sectors with clear regional interest—such as financial services (where regulatory cooperation already exists through central bank networks) or healthcare (where cross-border medical AI applications raise shared governance challenges). Successful sectoral pilots build institutional learning and political coalitions for broader application.

Leveraging existing institutions. ASEAN possesses established mechanisms—Senior Labour Officials Meeting (SLOM), ASEAN Finance Ministers' Meeting (AFMM), ASEAN Telecommunications and Information Technology Ministers Meeting (TELMIN)—that could incorporate velocity governance dialogues within existing mandates rather than requiring new bureaucratic structures. This reduces institutional startup costs and embeds coordination in proven forums.

This modular pathway acknowledges that binding, comprehensive ASEAN-wide velocity governance is politically premature. However, graduated coordination—beginning with soft mechanisms and scaling toward harder integration as capacity and trust develop—aligns with ASEAN's established integration patterns while addressing regulatory arbitrage risks (Acharya, 2009). The objective is not perfect harmonization, but sufficient coordination to prevent destabilizing races to the bottom while respecting developmental diversity.

8.5 Intergenerational Stakes and Upward Mobility

For emerging economies with youthful labor forces, entry-level bottlenecks in decision-layer occupations may carry disproportionate political weight. Upward mobility expectations often anchor social trust and political legitimacy (Lipset, 1959). Youth cohorts facing structural compression in analytical, managerial, and professional entry pathways may experience prolonged underemployment, skill underutilization, and blocked advancement—outcomes that generate political alienation even when headline unemployment remains low.

Malaysia's persistent youth unemployment ratio (approximately 3–3.5× the national rate, as documented in Section 9) illustrates this dynamic. The scale of regional exposure is considerable. The AI Asia Pacific Institute (2025), in collaboration with LinkedIn, estimates that 164 million ASEAN workers — over half the region's labor force — face AI-driven disruption, with skills requirements for jobs in Southeast Asia projected to change by 72 percent between 2016 and 2030, compared to 40 percent in the preceding nine years. Within Malaysia specifically, TalentCorp (2025) documented displacement risk across ten key economic sectors in an official multi-volume study commissioned by the Ministry of Human Resources — with aggregate parliamentary reporting placing the at-risk figure at approximately 620,000 jobs, concentrated in service and administrative roles that have historically served as entry-level pathways into formal middle-income employment. These are not marginal occupations; they are the rungs of the upward mobility ladder. Institutional preemption in such contexts is not merely economic prudence; it is a safeguard against intergenerational dislocation that could undermine social cohesion and political stability. Countries with aging populations and declining labor force participation may experience compression differently than those with youth bulges and rising educational attainment. ASEAN's developmental heterogeneity thus creates variance in both compression vulnerability and absorption urgency.

8.6 Stability as Competitive Advantage

Stability, in open economies, is itself a competitive asset. Foreign investors value predictable regulatory environments, sustainable labor markets, and institutional coherence alongside cost advantages and market access. By aligning acceleration with absorption through velocity governance, states may avoid the institutional fragility, social instability, and political unpredictability that undermine long-term investment attractiveness.

Countries that implement institutional preemption position themselves as stable, forward-looking jurisdictions rather than regulatory laggards or unstable environments. Singapore's reputation for regulatory sophistication and institutional quality attracts investment despite higher costs relative to neighbors; Malaysia could similarly leverage velocity-sensitive governance as marker of institutional maturity and long-term stability. This reframes institutional preemption from "competitiveness cost" to "institutional differentiation

strategy"—a positioning particularly valuable for middle-income economies seeking to move up value chains (Rodrik, 2011).

The regional application challenge is thus threefold: designing two-speed models that preserve competitiveness while protecting legitimacy; building modular ASEAN coordination that addresses arbitrage without triggering sovereignty resistance; and positioning velocity governance as stability-enhancing rather than growth-constraining. Open mid-sized economies cannot unilaterally control AI development trajectories, but they can control domestic deployment velocity and regional coordination mechanisms—precisely the domains where institutional preemption offers greatest leverage.

8.7 Developmental Heterogeneity and Absorptive Variance

The Compression–Absorption Framework is developmentally neutral in structure but developmentally variable in application. The stability condition $C(t) \leq A(t)$ does not assume uniform absorptive capacity across jurisdictions. Institutional thickness, fiscal depth, demographic composition, labor market formalization, and administrative capability vary substantially across advanced, aging, and middle-income economies. These differences shape both baseline absorptive capacity A_0 and the rate at which it can expand.

Advanced high-capacity economies—such as the United States, core EU states, and Singapore—possess mature welfare systems, diversified tax bases, established retraining institutions, and strong regulatory infrastructure. In these contexts, compression pressures are less likely to manifest first as unemployment shocks and more likely to appear through wage polarization, professional hollowing, and authority concentration within decision-layer occupations. The governance risk centers on inequality management and gradual erosion of corrective optionality rather than abrupt systemic destabilization.

Aging advanced economies—such as Japan and South Korea—exhibit high institutional capacity but face demographic contraction and labor scarcity. In these contexts, automation may function as partial substitution for declining workforce supply rather than pure displacement. Compression dynamics therefore differ by sector: labor-short industries may experience stabilization effects, while decision-layer automation still raises authority and institutional concentration concerns. The stability question remains velocity-relative, but demographic structure alters compression manifestation.

Middle-income and developing economies—including Malaysia and several ASEAN peers—often operate with thinner fiscal buffers, narrower tax bases, youth-heavy labor forces, and evolving retraining systems. Where upward mobility expectations anchor political legitimacy, entry-level bottlenecks and stalled professional pathways may generate disproportionate social strain. In such contexts, even moderate compression velocity may exceed absorptive capacity more rapidly, particularly where development models depend on labor-intensive exports or service outsourcing vulnerable to automation substitution.

Demographic composition further conditions absorption dynamics. Aging societies may tolerate higher automation velocity in labor-short sectors, whereas youth-bulge economies must continuously generate new entry pathways to sustain mobility expectations. The same deployment velocity may therefore produce distinct political economy effects across jurisdictions.

A further structural feature requiring explicit treatment is the informal-sector buffering mechanism. In advanced economies, displacement from formal employment generates observable signals — unemployment claims, benefit enrollment, wage compression in remaining roles — that detection systems can monitor and that trigger policy responses. In ASEAN labor markets, this signal pathway is structurally attenuated. Displaced formal workers often absorb into informal employment: self-employment, gig arrangements, informal trade, and subsistence agriculture. The displacement is real and the welfare loss is substantial, but the aggregate unemployment rate may remain stable. Arias et al. (2025) document this dynamic empirically: across five ASEAN countries, robot adoption 2018–2022 displaced approximately 1.4 million low-skilled formal workers, the majority of whom transitioned to informal sector employment rather than unemployment. Causal estimates showed lower average wages for the least-educated workers in higher-exposure locations, confirming welfare deterioration despite stable headline employment.

This creates a formal measurement requirement. In developing-economy contexts, $G(t)$ must be evaluated against a formal-sector-adjusted $A(t)$ that excludes informal absorption. Including informal employment as absorbed capacity artificially inflates $A(t)$ and produces false-negative $G(t)$ readings—the stability condition appears satisfied when structural displacement is in fact occurring. A conservative approach counts only formal-sector redeployment toward $A(t)$: documented cross-sector reemployment with social insurance re-enrollment. Informal absorption is tracked separately as a welfare-deterioration indicator. Without this adjustment, the framework’s detection architecture will systematically underestimate $G(t)$ in the economies where anticipatory governance is most needed.

The governance implication is that standard $G(t)$ detection architectures — designed around unemployment rate monitoring and formal sector wage data — will systematically underdetect compression in ASEAN contexts. $A(t)$ will appear artificially high because informal sector absorption masks displacement from formal roles. For the Compression–Absorption Framework to function as a governance instrument in developing economies, its Detection Layer (Section 6.1) must incorporate informal sector indicators: informal employment share by sector, earnings dispersion between formal and informal segments, transitions from formal to informal employment tracked through social insurance de-enrollment, and gig platform participation rates. These data are more difficult to collect than formal sector statistics, but they are the signals that developing-economy compression generates. Section 9.6 addresses data infrastructure requirements for Malaysia with this constraint in mind.

These differences do not alter the core framework; they reinforce it. Velocity governance must be calibrated to jurisdiction-specific absorptive baselines rather than applied uniformly. Formally, this implies heterogeneity in baseline A_0 and scaling elasticity of absorptive capacity across jurisdictions — parameters that must be estimated empirically for each context before threshold-triggered governance responses are calibrated.

Table 8.1

Developmental and Demographic Variation in Compression–Absorption Dynamics

Dimension	Advanced High-Capacity Economies (US, EU Core, Singapore)	Aging Advanced Economies (Japan, South Korea)	Middle-Income / Developing Economies (Malaysia, ASEAN peers)
Baseline Absorptive Capacity (A ₀)	High: mature welfare systems, fiscal depth, retraining infrastructure	High institutional capacity but constrained by demographic contraction	Moderate to low: thinner fiscal buffers, evolving retraining systems
Demographic Profile	Aging but stabilized populations	Rapid aging; shrinking labor force; low fertility	Youth-heavy or expanding labor force
Labor Market Structure	Highly formalized; large professional decision-layer sectors	High formalization; sectoral labor shortages	Mixed formal/informal sectors; growing but uneven professionalization
Primary Compression Manifestation	Wage polarization; middle-tier professional hollowing; authority concentration	Substitution for labor scarcity; selective professional compression	Entry-level bottlenecks; stalled upward mobility; youth underemployment
Immediate Political Sensitivity	Inequality and middle-class erosion	Pension sustainability; productivity stabilization	Intergenerational mobility expectations; employment legitimacy
Fiscal Shock Absorption	Strong redistributive capacity	Strong but pressured by aging-related fiscal burdens	More limited fiscal buffers; reliance on targeted transfers
Lock-In Risk Profile	Gradual entrenchment via capital deepening and institutional fatigue	Lower in labor-short sectors; persistent in decision-layer automation	Faster legitimacy strain if compression narrows entry pathways
Velocity Threshold Tolerance	Moderate to high tolerance before systemic strain	Higher tolerance in labor-scarce sectors; lower in governance-critical domains	Lower tolerance due to thinner absorptive margins
Governance Priority	Managing inequality and preserving distributed authority	Balancing automation with demographic stabilization	Protecting mobility ladders and scaling absorptive institutions

The comparative patterns summarized above demonstrate that baseline absorptive capacity A_0 and its endogenous scaling differ structurally across developmental and demographic contexts; accordingly, the stability condition $C(t) \leq A(t)$ must be evaluated relative to jurisdiction-specific institutional thickness and demographic composition rather than assumed uniform across states.

Institutional preemption is therefore not a one-size-fits-all doctrine. It is a velocity-sensitive stability condition whose operational thresholds depend on local absorptive margins. The relevant question is not whether automation is beneficial in the abstract, but whether deployment velocity remains compatible with institution-specific adaptation capacity.

8.8 Structural Early Warning and Trigger Architecture

8.8.1 From Stability Condition to Preventive Governance Logic

The compression–absorption condition developed earlier functions not only as an explanatory account of institutional lag, but as a preventive governance logic. If structural instability risk increases when compression velocity persistently exceeds absorptive capacity, then the relevant governance problem becomes anticipatory rather than retrospective: can widening asymmetry be detected before reversibility erodes?

This section translates the velocity–absorption framework into a structured early-warning architecture. The argument remains conditional and non-alarmist. Intervention is not triggered by technological novelty, sectoral disruption, or isolated displacement. It is triggered only where sustained asymmetry becomes observable across multiple domains.

The objective is neither precautionary inhibition nor technological deceleration per se, but the preservation of institutional optionality under accelerating compression.

8.8.2 Structural Early Warning Domains

The stability condition implies four analytically distinct but interrelated domains through which asymmetry may manifest.

(1) Accelerated Functional Compression

Compression becomes structurally salient when automation deployment within labour-dense or decision-critical functions exhibits patterned acceleration across firms and sectors without commensurate adaptation pathways. The relevant signal is not individual role transformation, but clustered diffusion within contiguous occupational or functional categories over compressed time intervals.

Structural risk arises where redesign tempo consistently exceeds redeployment capacity, producing cumulative rather than transitional displacement pressure.

(2) Institutional Response Lag

Absorptive capacity depends upon regulatory, administrative, and oversight adaptation. Institutional lag becomes destabilising when governance response cycles systematically trail deployment cycles across multiple iterations. Reactive legislative updates, insufficient technical audit capacity, and the absence of structured deployment transparency mechanisms may signal cumulative response deficit.

The critical variable is not delay in isolation, but delay under conditions of accelerating compression.

(3) Reversibility Erosion

The most consequential dimension of asymmetry concerns reversibility. Structural fragility increases where rollback capacity declines due to infrastructural embedding, vendor lock-in, deep system integration, or prohibitive decommissioning costs. When algorithmic infrastructures become non-modular components of core administrative or market functions, corrective optionality narrows.

Compression under conditions of reversibility erosion transforms adaptive pressure into structural constraint.

(4) Authority Concentration

Compression also reallocates decision authority. Structural risk intensifies where operational discretion migrates from distributed institutional actors toward centralised algorithmic systems without proportional accountability architecture. Indicators may include diminished feasibility of meaningful human override, opaque optimisation logic, and concentration of infrastructural control within limited vendor ecosystems.

The concern is not automation itself, but asymmetrical authority consolidation under high-velocity deployment.

8.8.3 Tiered Trigger Logic

Early warning requires proportional activation criteria. A tiered trigger logic preserves analytical clarity while avoiding blunt intervention mechanisms.

Tier I: Monitoring.

Where indicators suggest rising compression but absorptive alignment remains intact, governance posture remains informational—enhanced transparency, deployment mapping, and institutional capacity reinforcement.

Tier II: Stability Risk.

Where sustained asymmetry becomes observable across multiple domains, modulatory responses become appropriate. These may include structured deployment review mechanisms, conditional approval gates within high-impact domains, and mandatory absorption impact assessments.

Tier III: Lock-In Risk.

Where reversibility erosion coincides with persistent asymmetry, temporary deceleration measures within narrowly defined domains may be warranted. Such measures must remain time-bound, reviewable, and proportionate to the stability risk identified.

The governing principle is velocity-based rather than technology-based. Activation depends on sustained asymmetry, not innovation intensity.

8.8.4 Proportionality, Calibration, and Preservation of Optionality

This framework does not presume technological harm nor advocate precautionary paralysis. It specifies a stability condition under which governance priority shifts from acceleration to alignment. Where compression persistently exceeds absorption and reversibility declines, institutional fragility increases irrespective of short-term efficiency gains.

The preservation of structural optionality—rather than the suppression of innovation—constitutes the governing objective.

The calibration of “sustained asymmetry” into actionable thresholds is necessarily jurisdiction-specific. Operationalising this tiered logic depends upon the development of the data infrastructure and reporting capacity discussed in Section 9.6. Without such informational architecture, warning domains remain conceptual. With it, they become empirically calibratable governance instruments capable of translating structural theory into proportionate activation criteria.

In this sense, preventive governance is not a brake on technological progress, but a mechanism for maintaining reversibility under conditions of accelerating compression.

9. Illustrative Application: Applying the Compression–Absorption Framework to Malaysia (2015–2025)

This section applies the Compression–Absorption Framework as an illustrative application rather than a definitive empirical test. Its purpose is methodological: to demonstrate how $C(t)$, $A(t)$, and $G(t)$ can be operationalized using publicly available data, and to identify what a fully specified empirical application would require. The Malaysian case is selected for illustrative tractability, not claimed as conclusive evidence of compression dynamics. No causal attribution of observed labor trends to AI-specific mechanisms is asserted; the available macro-level data does not yet permit such attribution. To anticipate a direct reviewer concern: Section 9 does not validate the framework. It does not present AI-specific deployment metrics. It uses macro labor market data that cannot distinguish AI-driven from broader automation effects. The data show historically low unemployment, which is consistent with the absorptive buffering hypothesis but not diagnostic of compression dynamics.

The honest characterization is that Section 9 is a data architecture exercise: it specifies what must be measured, demonstrates that existing data sources can partially operationalize the framework's components, and identifies the gaps that future data collection must close. Empirical validation of the framework requires AI-specific deployment metrics linked to occupation-level task substitution data — data that do not yet exist at the required granularity. Presenting this as anything more would be methodologically dishonest. To state this plainly: this section does not claim that Malaysia is already being compressed by AI — it demonstrates that the framework is operationalizable under current data constraints. The section's length relative to the theoretical core reflects a deliberate methodological choice: anticipatory governance frameworks risk remaining speculative unless researchers and agencies can see precisely what data, at what disaggregation, operationalize each component. Sections 9.6 and 9.7 serve this function and could in principle be developed into a companion policy paper; their inclusion here is justified by the argument that the gap between framework design and empirical implementation constitutes part of the contribution for frameworks in this class. What the exercise does establish is that the framework's diagnostic logic is applicable to real institutional contexts, and that observable indicators are directionally consistent with the compression pressures the theory predicts.

A temporal note applies to this section. All Malaysian data cited here — DOSM Labour Force Statistics, DOSM Employee Wages Statistics, MOF budget allocations, PERKESO disbursement figures, and the Gig Workers Act 2025 — are drawn from published official sources available as of early 2026 and treated as contemporaneous observations within the study window. They are not forward projections or modelled scenarios. Readers applying this framework prospectively — for instance, constructing $C(t)$ estimates for 2027 or 2028 — would necessarily be working with projected or scenario-based values for some components.

The framework’s design accommodates this: it is intended to be operationalizable in advance of crisis-level data, using the best available indicators at the time of application. Where this illustrative application uses verified observations, future applications may substitute scenario parameters at the same structural positions without altering the diagnostic logic.

This empirical illustration relies exclusively on official releases from the Department of Statistics Malaysia (DOSM), Ministry of Finance (MOF), and related government communications (Department of Statistics Malaysia [DOSM], 2023a, 2024a, 2025a; Ministry of Finance Malaysia [MOF], 2025). Where full time-series data through 2025 are unavailable, the analysis clearly identifies constraints.

9.1 Measurement Methodology

Consistent with Appendix A, this section distinguishes between compression proxies (C-proxies) and absorption proxies (A-proxies).

Compression Proxies

1. **National unemployment rate** – Monthly and annual labor force statistics (Department of Statistics Malaysia [DOSM], 2025a).
2. **Youth unemployment rate (15–24)** – Annual Labour Force Survey releases (DOSM, 2024b).
3. **Youth-to-national unemployment ratio** – Calculated from official rates (DOSM, 2024b, 2025a).
4. **Long-term unemployment share** – Percentage of unemployed individuals jobless for more than one year (DOSM, 2025a).
5. **Wage dispersion (90th/10th percentile ratio)** – Formal sector wage statistics (DOSM, 2025b).
6. **Compensation of employees (CE) share of GDP** – Income Approach to GDP releases (DOSM, 2023a, 2024a).

These variables do not measure AI deployment directly. Rather, they capture macro-level labor compression conditions—such as task substitution, wage polarization, and participation strain—consistent with established automation literature (Acemoglu & Restrepo, 2018; Autor, 2013).

Absorption Proxies

1. **Cash transfer expansion (STR/SARA allocations)** – Federal budget allocations and recipient coverage (Ministry of Finance Malaysia [MOF], 2025).
2. **PERKESO benefit payouts and recipient coverage** – Official annual disbursement statistics and coverage figures (Social Security Organisation [PERKESO], 2025).
3. **Gig Workers Act 2025 (Act 872)** – Statutory formalization of gig worker protections (Malaysia, 2025).

4. **Expanded social protection initiatives (e.g., LINDUNG 24 Jam)** – Official policy expansion statements and implementation details (Social Security Organisation [PERKESO], 2025; Ministry of Human Resources Malaysia [MOHR], 2025).

These indicators reflect short-run buffering capacity and institutional responsiveness. They do not, by themselves, measure structural labor redeployment or long-term occupational reallocation.

9.2 Core Labor Market Indicators (2023–2025)

Table 9.1. Verified Compression Indicators

Indicator	2023	2024	2025
Compensation of Employees Share of GDP	33.1% (DOSM, 2023a)	33.6% (DOSM, 2024a)	Not yet published
National Unemployment Rate	3.4% (annual context, DOSM, 2024a)	3.2% (annual, DOSM, 2024a)	2.9% (Nov 2025, DOSM, 2025a)
Youth Unemployment (15–24)	11.0% (DOSM, 2024b)	10.3% (DOSM, 2024b)	10.1% (DOSM, 2025a)
Long-Term Unemployment Share	—	5.8% (DOSM, 2024b)	5.0% (DOSM, 2025a)
Wage 90/10 Ratio (Formal Sector)	—	—	≈5.5× (DOSM, 2025b)

Observations

1. **Headline unemployment is historically low** (2.9% in late 2025), suggesting a tight labor market (DOSM, 2025a).
2. **Youth unemployment remains structurally elevated** (~10%), even as it modestly improves (DOSM, 2024b, 2025a).
3. **Wage dispersion is substantial** (≈5.5× between 90th and 10th percentile earners), indicating persistent inequality within formal employment (DOSM, 2025b).
4. **CE share rose from 2023 to 2024**, contradicting a simplistic narrative of continuous labor share decline (DOSM, 2023a, 2024a).

Taken together, the data suggest that while aggregate employment remains strong, distributional and entry-level pressures persist.

9.3 Absorptive Capacity Indicators (2024–2025)

Table 9.2. Verified Absorption Indicators

Indicator	2024	2025
STR + SARA Allocation	RM10b baseline (Ministry of Finance Malaysia [MOF], 2025 context)	RM13b; ~9 million recipients (MOF, 2025)
PERKESO Benefit Payouts	—	RM6.65b; >1.15 million recipients (Social Security Organisation [PERKESO], 2025)
Legal Reform	—	Gig Workers Act 2025 (Malaysia, 2025)

Observations

Malaysia expanded fiscal support significantly in 2025, increasing direct household transfers and strengthening social insurance coverage (MOF, 2025; Social Security Organisation [PERKESO], 2025). These measures reflect a measurable expansion of short-term absorptive buffering capacity within the national labor and welfare system.

However, fiscal stabilization differs conceptually from structural absorption. While transfer expansion and social insurance payouts mitigate income shocks, they do not by themselves demonstrate long-term labor redeployment, occupational mobility, or skills upgrading. Structural absorption requires evidence of sustained retraining pipelines, sectoral reallocation, and productivity-linked employment transitions beyond income buffering mechanisms.

9.4 Governance Gap Interpretation

The exercise deliberately exposes data constraints: task-level substitution rates, sectoral automation penetration coefficients, and occupational entry velocity data are not yet available at the granularity required for full $C(t)$ and $A(t)$ quantification from public Malaysian sources. These gaps are documented as part of the illustrative application’s contribution—identifying precisely what data collection is necessary for the framework’s empirical operationalization. Where available indicators permit directional scoring, results are presented as illustrative estimates. Where they do not, the section specifies the required data architecture. Pending complete annual series through 2025, precise numeric calculation of $C(t)$, $A(t)$, and $G(t)$ remains provisional.

What the verified data allow:

- Compression pressures remain present in distributional indicators (youth unemployment; wage dispersion).
- Absorptive capacity expanded materially in 2025 through fiscal and legislative intervention.
- No macro-level unemployment shock is observable.

Thus, Malaysia’s short-run governance gap appears contained, though structural redeployment capacity remains an open empirical question.

Under the Compression–Absorption Framework, this pattern is consistent with **effective buffering combined with persistent structural strain**.

9.5 What This Evidence Demonstrates

Supported by official statistics:

- Historically low unemployment (DOSM, 2025a).
- Persistent youth employment gap (~3.4× national rate).
- Significant wage dispersion (DOSM, 2025b).
- Large-scale fiscal expansion and institutional reform (MOF, 2025; Government of Malaysia, 2025).

Not demonstrated:

- AI-specific causality.
 - Permanent closure of the governance gap.
 - Long-term redeployment sufficiency.
-

9.6 Data Collection Infrastructure: Actionable Specification

The framework’s anticipatory design requires data infrastructure to be built before crisis-level compression is observable. The following specification moves from “these data do not yet exist” to “here is exactly what statistical agencies should begin collecting in 2025–2026 to enable $C(t)$ measurement when needed.”

Priority 1: Immediate Collection (Existing Infrastructure, Enhanced Disaggregation)

These datasets are producible within current statistical agency capacity, requiring enhanced disaggregation rather than new survey instruments. DOSM should extend the Labour Force Survey to include occupational task composition breakdowns by 2-digit ISCO code, enabling task substitution rate tracking across years. The Salaries and Wages Survey should be published annually (currently biennial) with consistent 10th, 25th, 50th, 75th, and 90th percentile series by sector, enabling wage polarization indices without interpolation. The Ministry of Higher Education Graduate Tracer Study should be disaggregated by occupational destination category — not merely employment status — to enable entry bottleneck rate calculation. EPF and PERKESO administrative records should be linked to produce cross-sector mobility tracking for workers who change employer and sector within 24-month windows. Critically, PERKESO de-enrollment data should be used to track formal-to-informal sector transitions — workers who exit formal employment without registering for benefits represent a key informal-sector absorption signal that standard unemployment statistics will not capture (consistent with the informal-sector buffering mechanism identified

in Section 8.7). These enhancements require administrative instruction and budget allocation, not new methodological development. Target timeline: 2025–2026 collection cycles.

Priority 2: New Data Collection (Requires New Survey Instruments)

Firm-level automation adoption surveys should be administered through DOSM’s existing Business Tendency Survey infrastructure, adding a module capturing: proportion of tasks performed by automated systems by function (administrative, analytical, decision-layer, physical); year of deployment for AI-specific tools; and human headcount changes attributable to automation rather than demand. HRD Corp levy-funded retraining programs should be tracked by outcome category: employed in same sector, employed in different sector, still retraining, not employed. This converts HRD Corp data from expenditure tracking into structural absorption measurement. MOHR should establish an occupational task database updated annually, classifying Malaysian Standard Occupation Codes by routine versus non-routine and manual versus cognitive dimensions, enabling direct mapping to Autor-style automation vulnerability scores. Target timeline: 2026–2027 first data release.

Priority 3: Methodological Development (Requires Research Partnerships)

Attribution models separating AI-specific deployment effects from broader automation trends require longitudinal firm-level panels linking technology adoption records to labour market outcomes. This requires partnership between DOSM, MDEC (which tracks digital adoption), and EPF (which tracks employment continuity). Threshold calibration — determining what $G(t)$ level triggers preemptive governance response — requires cross-country comparison using historical automation episodes, building on Acemoglu and Restrepo (2020)’s robot exposure methodology applied to AI-specific tools. Target timeline: 2027–2029 calibration studies.

Lead–Lag Feasibility Assessment: Preemption Windows by Indicator

Not all $C(t)$ and $A(t)$ components are equally timely. Understanding the lead–lag structure of each indicator is essential for determining realistic preemption windows — the interval between observable signal and governance response. The following assessment classifies each component by measurement lag under current Malaysian statistical infrastructure.

Real-time indicators (0–3 month lag): National and youth unemployment rates are published monthly by DOSM with a 4–6 week lag. These are the earliest available signals of compression but are the least specific — headline unemployment is a lagging indicator even within this category, as it reflects workforce exits already completed rather than compression in progress. EPF contribution volumes by sector provide a near-real-time proxy for formal employment levels, with monthly lag. These indicators provide the earliest warning but require supplementation from slower, more specific measures.

Short-lag indicators (3–6 month lag): Wage data from DOSM’s quarterly Salaries and Wages releases (once annualised publication is converted to quarterly) provide wage polarization signals within one to two quarters of the relevant period. PERKESO benefit claims provide a 3–4 month lag proxy for involuntary displacement, as claimants must register within 30 days of job loss. HRD Corp retraining enrollment figures are available within one quarter and provide an early $A(t)$ signal.

Medium-lag indicators (6–18 month lag): Graduate tracer data lags by 6–12 months as it tracks employment status approximately six months after graduation. Occupational entry rate data derived from EPF records requires longitudinal matching and is typically available with a 12–18 month lag. Firm-level automation adoption survey data (once established) would be available annually with a 12–18 month lag due to survey design, fieldwork, and processing. This category contains the most diagnostically specific $C(t)$ components — task substitution rates, decision-layer entry rate decline — but they are the slowest to arrive.

Implication for preemption windows: The lead–lag structure implies a structural preemption constraint. The most specific compression indicators lag by 6–18 months; preemptive governance responses require political deliberation, regulatory drafting, and implementation periods of 6–24 months. The total cycle from measurable compression signal to effective institutional response is therefore 12–42 months under current conditions. For the preemption doctrine to be operationally credible, governance frameworks must be designed and threshold parameters set in advance, so that when medium-lag indicators confirm elevated $C(t)$, the response architecture is already in place and can be activated administratively rather than requiring fresh legislative processes. This reinforces the core argument for institutional preemption: the monitoring and response infrastructure must precede the crisis it is designed to detect.

9.7 Counter-Scenario Analysis: Framework Applicability Across Divergent Futures

A diagnostic framework must be discriminating: it should identify elevated risk when conditions warrant, and withhold alarm when they do not. This section applies the Compression–Absorption Framework to two alternative deployment scenarios alongside the baseline case, demonstrating that the framework produces differentiated outputs across futures rather than a single predetermined conclusion.

Scenario A (Baseline): Sustained High-Velocity Deployment

Assumptions: AI deployment in analytical and decision-layer functions continues at current trajectory; task substitution rates in financial services, logistics, and professional services reach 8–12% per annum by 2027–2029; retraining absorption remains constrained by current program capacity and fiscal allocation. Under these conditions, $C(t)$ rises from the illustrative 0.71 baseline (Appendix A4) toward 0.80–0.85 within three years. $A(t)$ expands incrementally — fiscal buffers increase, Gig Workers Act provisions mature — but absorption velocity lags compression by a sustained margin. $G(t)$ remains positive and widens. Framework diagnosis: elevated lock-in risk; threshold-triggered review mechanisms and phased rollout conditions are activated under the Section 8 governance architecture. Institutional preemption is warranted.

Scenario B: Deployment Plateau with Moderate Absorption

Assumptions: AI capability diffusion slows after 2026 due to regulatory friction, compute constraints, or market saturation in early-adopter sectors; task substitution rates stabilize at 3–4% per annum; concurrent investment in retraining infrastructure (HRD Corp expansion, graduate tracer integration, sectoral reskilling pipelines) raises retraining velocity from 35% to 55% redeployment within 18 months. Under these conditions, $C(t)$ stabilizes around 0.45–0.50. $A(t)$ rises toward 0.48–0.55 as absorption infrastructure matures. $G(t)$ approaches zero or turns marginally negative, indicating absorption keeping pace with compression. Framework diagnosis: stability condition approached; monitoring continues but preemptive velocity controls are not triggered. Institutional preemption remains on standby — detection infrastructure is maintained precisely because this scenario could shift toward Scenario A if deployment resumes. The framework correctly produces a non-alarm diagnosis without being switched off.

Scenario C: Accelerated Deployment with Weak Absorption

Assumptions: Generative AI adoption accelerates beyond current projections; decision-layer automation penetrates legal, medical, and educational functions at scale by 2027; political economy constraints delay retraining reform (Section 8 political economy analysis); fiscal capacity remains constrained by debt service obligations. Under these conditions, $C(t)$ rises rapidly toward 0.90+. $A(t)$ remains depressed at 0.25–0.30 as retraining systems are overwhelmed and mobility mechanisms stall. $G(t)$ reaches 0.60–0.65, indicating severe compression–absorption asymmetry. Cumulative lock-in pressure $P(T)$ accumulates rapidly. Framework diagnosis: critical threshold breached; emergency institutional preemption protocols, including mandatory phased deployment restrictions in highest-compression sectors and emergency fiscal reallocation, are indicated.

The three scenarios demonstrate that the Compression–Absorption Framework is not calibrated to produce a single predetermined warning. It is sensitive to the actual relationship between $C(t)$ and $A(t)$: it withholds preemptive signals when absorption keeps pace (Scenario B), escalates signals proportionally to gap magnitude (Scenarios A and C), and provides a graduated response architecture rather than a binary alarm. This diagnostic discrimination is the framework’s central methodological contribution — it enables proportionate governance across futures rather than locking policymakers into a single anticipatory posture.

9.8 Summary

Malaysia’s 2024–2025 data suggest substantial short-run absorptive buffering alongside persistent structural labor pressures. This supports a moderated interpretation of the framework: institutional preemption is not an emergency brake, but a velocity governance mechanism ensuring that structural compression does not outpace redeployment capacity.

10. Conclusion: Governing Acceleration Before Entrenchment

Labor-compressing AI systems introduce a structural governance challenge centered on velocity. The question is not whether technological transformation will occur, but whether institutional adaptation can keep pace. When compression velocity persistently exceeds absorptive capacity, structural drift may advance before political systems recognize its cumulative effect.

This paper has proposed the Compression–Absorption Framework as a stability condition for labor-compressing AI. It has argued that sustained mismatches between compression and adaptation generate cumulative lock-in pressure, narrowing corrective optionality. The doctrine of Institutional Preemption reframes AI governance as anticipatory equilibrium management rather than reactive stabilization. It calls for detection mechanisms, velocity modulation, strengthened absorptive capacity, preserved accountable authority, and regional coordination in open economies.

The absence of crisis does not invalidate preemption. Preventive governance often manifests as uneventful stability. If severe disruption does not materialize, it may be because compression never exceeded absorptive capacity—or because institutional safeguards aligned acceleration with adaptation.

The framework's value lies in its applicability across divergent futures. Whether labor-compressing AI capabilities advance gradually along current trajectories, plateau below transformative thresholds, or accelerate beyond present expectations, velocity-sensitive governance infrastructure provides institutional resilience. Uncertainty about deployment speed and capability trajectories strengthens rather than weakens the case for early preparation. The cost of building monitoring systems, establishing velocity thresholds, and preserving corrective optionality remains modest compared to the cost of institutional reconstruction once structural compression has hardened into constitutional constraint.

AI governance must therefore evolve beyond safety and fairness alone. It must incorporate macro-institutional velocity as a central variable. The preservation of human agency depends less on prohibiting machines and more on ensuring that technological acceleration remains compatible with institutional equilibrium—regardless of whether that acceleration proceeds incrementally or intensifies unpredictably.

The central policy challenge is not transformation, but timing. Governance must adapt quickly enough to preserve human agency and institutional reversibility before velocity asymmetry converts structural compression into irreversible fact. The Compression–Absorption Framework enables states to measure, monitor, and modulate that velocity before the window for correction closes. This is the essence of institutional preemption: acting when intervention remains feasible, rather than reacting when entrenchment has already occurred.

References

- Acharya, A. (2009). *Constructing a security community in Southeast Asia: ASEAN and the problem of regional order* (2nd ed.). Routledge.
- Acemoglu, D., & Restrepo, P. (2018). Artificial intelligence, automation, and work. In A. Agrawal, J. Gans, & A. Goldfarb (Eds.), *The economics of artificial intelligence: An agenda* (pp. 197–236). University of Chicago Press.
- Acemoglu, D., & Restrepo, P. (2019). Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives*, 33(2), 3–30.
<https://doi.org/10.1257/jep.33.2.3>
- Acemoglu, D., & Restrepo, P. (2020). Robots and jobs: Evidence from U.S. labor markets. *Journal of Political Economy*, 128(6), 2188–2244. <https://doi.org/10.1086/705716>
- Admati, A., & Hellwig, M. (2013). *The bankers' new clothes: What's wrong with banking and what to do about it*. Princeton University Press.
- AI Asia Pacific Institute. (2025). AI and workforce transformation: Unlocking pathways to inclusive growth in ASEAN developing economies. AI Asia Pacific Institute & LinkedIn.
<https://aiasiapacific.org/our-work/2025-ai-and-workforce-transformation-unlocking-pathways-to-inclusive-growth-in-asean-developing-economies/>
- Algan, Y., Guriev, S., Papaioannou, E., & Passari, E. (2017). The European trust crisis and the rise of populism. *Brookings Papers on Economic Activity*, 2017(2), 309–400.
<https://doi.org/10.1353/eca.2017.0015>
- Arias, O., Fukuzawa, D., Le, D. T., & Mattoo, A. (2025). Future jobs: Robots, artificial intelligence, and digital platforms in East Asia and Pacific. World Bank.
<https://doi.org/10.1596/978-1-4648-2191-9>
- Arthur, W. B. (1989). Competing technologies, increasing returns, and lock-in by historical events. *The Economic Journal*, 99(394), 116–131. <https://doi.org/10.2307/2234208>
- Autor, D. H. (2013). The “task approach” to labor markets: An overview. *Journal for Labour Market Research*, 46(3), 185–199. <https://doi.org/10.1007/s12651-013-0128-z>
- Autor, D. H., & Dorn, D. (2013). The growth of low-skill service jobs and the polarization of the U.S. labor market. *American Economic Review*, 103(5), 1553–1597.
<https://doi.org/10.1257/aer.103.5.1553>
- Autor, D. H., Levy, F., & Murnane, R. J. (2003). The skill content of recent technological change: An empirical exploration. *Quarterly Journal of Economics*, 118(4), 1279–1333.
<https://doi.org/10.1162/003355303322552801>
- Beramendi, P., & Rueda, D. (2007). Social democracy constrained: Indirect taxation in industrialized democracies. *British Journal of Political Science*, 37(4), 619–641.
<https://doi.org/10.1017/S0007123407000341>

- Blanchard, O. J., & Summers, L. H. (1986). Hysteresis and the European unemployment problem. *NBER Macroeconomics Annual*, 1, 15–78. <https://doi.org/10.1086/654026>
- Bovens, M., & Zouridis, S. (2002). From street-level to system-level bureaucracies: How information and communication technology is transforming administrative discretion and constitutional control. *Public Administration Review*, 62(2), 174–184. <https://doi.org/10.1111/0033-3352.00168>
- Brynjolfsson, E., & McAfee, A. (2014). *The second machine age: Work, progress, and prosperity in a time of brilliant technologies*. W. W. Norton.
- Brynjolfsson, E., Rock, D., & Syverson, C. (2021). The productivity J-curve: How intangibles complement general purpose technologies. *American Economic Journal: Macroeconomics*, 13(1), 333–372. <https://doi.org/10.1257/mac.20180386>
- Caballero-Anthony, M., & Gong, L. (Eds.). (2020). Non-traditional security issues in ASEAN: Agendas for action. ISEAS–Yusof Ishak Institute.
- Citron, D. K., & Pasquale, F. (2014). The scored society: Due process for automated predictions. *Washington Law Review*, 89(1), 1–33.
- Crawford, K. (2021). *Atlas of AI: Power, politics, and the planetary costs of artificial intelligence*. Yale University Press.
- David, P. A. (1985). Clio and the economics of QWERTY. *The American Economic Review*, 75(2), 332–337.
- Davis, S. J., & von Wachter, T. (2011). Recessions and the costs of job loss. *Brookings Papers on Economic Activity*, 2011(2), 1–72. <https://doi.org/10.1353/eca.2011.0005>
- Department of Statistics Malaysia. (2023a). *Gross domestic product (income approach), 2023*.
- Department of Statistics Malaysia. (2024a). *Gross domestic product (income approach), 2024*.
- Department of Statistics Malaysia. (2024b). *Labour force survey report, 2024*.
- Department of Statistics Malaysia. (2025a). *Labour force statistics, November 2025*. Retrieved February 15, 2026, from <https://open.dosm.gov.my/dashboard/labour-market>
- Department of Statistics Malaysia. (2025b). *Employee wages statistics, formal sector, second quarter 2025*.
- Eubanks, V. (2018). *Automating inequality: How high-tech tools profile, police, and punish the poor*. St. Martin's Press.
- European Parliament & Council of the European Union. (2024). Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised

rules on artificial intelligence (Artificial Intelligence Act). Official Journal of the European Union, L 2024/1689. https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=OJ:L_202401689

Fukuyama, F. (2014). *Political order and political decay: From the industrial revolution to the globalization of democracy*. Farrar, Straus and Giroux.

Goldin, C., & Katz, L. F. (2008). *The race between education and technology*. Belknap Press of Harvard University Press.

Guston, D. H. (2014). Understanding ‘anticipatory governance’. *Social Studies of Science*, 44(2), 218–242. <https://doi.org/10.1177/0306312713508669>

Hacker, J. S. (2004). Privatizing risk without privatizing the welfare state: The hidden politics of social policy retrenchment in the United States. *American Political Science Review*, 98(2), 243–260. <https://doi.org/10.1017/S0003055404001121>

International Labour Organization. (2012). International standard classification of occupations: ISCO-08 (Vol. 1). ILO. https://www.ilo.org/wcmsp5/groups/public/---dgreports/---dcomm/---publ/documents/publication/wcms_172572.pdf

International Monetary Fund. (2024). Fiscal monitor, October 2024: Putting a lid on public debt. IMF. <https://www.imf.org/en/Publications/FM>

International Organization for Standardization & International Electrotechnical Commission. (2023). ISO/IEC 42001:2023: Information technology — Artificial intelligence — Management system. ISO. <https://www.iso.org/standard/81230.html>

Korinek, A., & Stiglitz, J. E. (2019). Artificial intelligence and its implications for income distribution and unemployment. In A. Agrawal, J. Gans, & A. Goldfarb (Eds.), *The economics of artificial intelligence: An agenda* (pp. 349–390). University of Chicago Press.

Kroft, K., Lange, F., & Notowidigdo, M. J. (2013). Duration dependence and labor market conditions: Evidence from a field experiment. *Quarterly Journal of Economics*, 128(3), 1123–1167. <https://doi.org/10.1093/qje/qjt015>

Lindert, P. H. (2004). *Growing public: Social spending and economic growth since the eighteenth century* (Vol. 1). Cambridge University Press.

Lipset, S. M. (1959). Some social requisites of democracy: Economic development and political legitimacy. *American Political Science Review*, 53(1), 69–105. <https://doi.org/10.2307/1951731>

Malaysia. (2025). *Gig Workers Act 2025 (Act 872)*. <https://jtksm.mohr.gov.my/sites/default/files/2026-01/Act%20872%20-%20GIG%20WORKERS%20ACT%202025.pdf>

Malaysia Digital Economy Corporation (MDEC). (2021). Malaysia national artificial intelligence roadmap 2021–2025. Ministry of Science, Technology and Innovation Malaysia. <https://mdec.my/ai/>

- Marchant, G. E., Allenby, B. R., & Herkert, J. R. (Eds.). (2008). *The growing gap between emerging technologies and legal-ethical oversight: The pacing problem*. Springer.
- Ministry of Finance Malaysia. (2025). *Distribution of RM13 billion STR/SARA aid*.
- Ministry of Human Resources Malaysia. (2025). *LINDUNG 24 Jam: Expanded social protection for gig workers*. Retrieved from <https://www.mohr.gov.my>
- Mokyr, J. (1990). *The lever of riches: Technological creativity and economic progress*. Oxford University Press.
- North, D. C. (1990). *Institutions, institutional change and economic performance*. Cambridge University Press.
- Ogburn, W. F. (1922). Social change with respect to culture and original nature. B. W. Huebsch.
- O’Neil, C. (2016). *Weapons of math destruction: How big data increases inequality and threatens democracy*. Crown.
- Organisation for Economic Co-operation and Development. (2019, updated 2023). Recommendation of the Council on Artificial Intelligence (OECD/LEGAL/0449). OECD Publishing. <https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0449>
- Pasquale, F. (2015). *The black box society: The secret algorithms that control money and information*. Harvard University Press.
- Pierson, P. (2000). Increasing returns, path dependence, and the study of politics. *American Political Science Review*, 94(2), 251–267. <https://doi.org/10.2307/2586011>
- Rodrik, D. (2011). *The globalization paradox: Democracy and the future of the world economy*. W. W. Norton.
- Schlozman, K. L., Verba, S., & Brady, H. E. (2012). *The unheavenly chorus: Unequal political voice and the broken promise of American democracy*. Princeton University Press.
- Social Security Organisation (PERKESO). (2025, October 17). *PERLINDUNGAN 24 JAM PERKESO memberi manfaat kepada lebih 9.6 juta pencarum* [Media release]. <https://www.perkeso.gov.my>
- Stigler, G. J. (1971). The theory of economic regulation. *The Bell Journal of Economics and Management Science*, 2(1), 3–21. <https://doi.org/10.2307/3003160>
- Sunstein, C. R. (2005). *Laws of fear: Beyond the precautionary principle*. Cambridge University Press.
- Talent Corporation Malaysia Berhad (TalentCorp). (2025). Impact study of artificial intelligence, digital, and green economy on the Malaysian workforce: Volume 2 — Sector: Wholesale and retail trade (N. Zulhaimi & A. S. Abdul Samat, Authors; N. Zulkifli, Ed.). Ministry of Human Resources Malaysia.

<https://www.talentcorp.com.my/images/uploads/publication/202/Impact-Study-of-Artificial-Intelligence-Digital-and-Green-Economy-on-the-Malaysian-Workforce-Volume-2-Sector-Wholesale-and-Retail-Trade-1747879709.pdf> [Aggregate sector-risk figure of ~620,000 jobs cited in Parliamentary Reply, Dewan Rakyat, 18 August 2025, drawing on full multi-volume series.]

Taleb, N. N. (2012). *Antifragile: Things that gain from disorder*. Random House.

Vogel, D. (1995). *Trading up: Consumer and environmental regulation in a global economy*. Harvard University Press.

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

Appendix A: Operationalizing the Compression–Absorption Framework

The Compression–Absorption Framework introduced in Section 3 formalizes a stability condition linking technological labor compression to institutional absorptive capacity. This appendix provides operational guidance for measurement, calibration, and empirical application. The objective is diagnostic clarity rather than predictive precision.

A1. Measuring Labor Compression $C(t)$

Labor compression should not be measured solely through headline unemployment rates. Task-based labor research demonstrates that automation often restructures occupations internally before manifesting as outright displacement (Autor, 2013; Acemoglu & Restrepo, 2018). Accordingly, $C(t)$ should be constructed as a composite indicator incorporating multiple dimensions.

Proposed Components of $C(t)$:

1. **Task Substitution Ratio**
Percentage of occupational tasks automated within affected sectors.
2. **Decision-Layer Automation Index**
Proportion of managerial, analytical, and supervisory functions executed by high-autonomy systems.
3. **Wage Compression Differential**
Divergence between productivity growth and median wage growth in automation-intensive sectors.
4. **Entry Bottleneck Rate**
Decline in first-time labor market absorption for decision-layer occupations.
5. **Labor Share Shift**
Change in national labor share of income relative to capital share.
6. **Occupational Persistence Decline**
Increase in duration of displacement or underemployment in affected sectors (Blanchard & Summers, 1986).

Composite $C(t)$ can be derived through weighted aggregation of these indicators, calibrated to national statistical capacity.

Figure A1 below provides a schematic representation of the framework’s core dynamics. The diagram illustrates three scenarios: (a) stable equilibrium in which $A(t)$ matches or exceeds $C(t)$; (b) an opening governance gap where $C(t)$ begins to exceed $A(t)$, generating positive $G(t)$; and (c) lock-in risk accumulation as $P(T)$ rises with sustained positive $G(t)$. The vertical axis represents the composite index values (0–1 normalized scale); the horizontal axis represents time. The preemption window is the interval between $G(t)$ first becoming positive and $P(T)$ crossing the estimated entrenchment threshold.

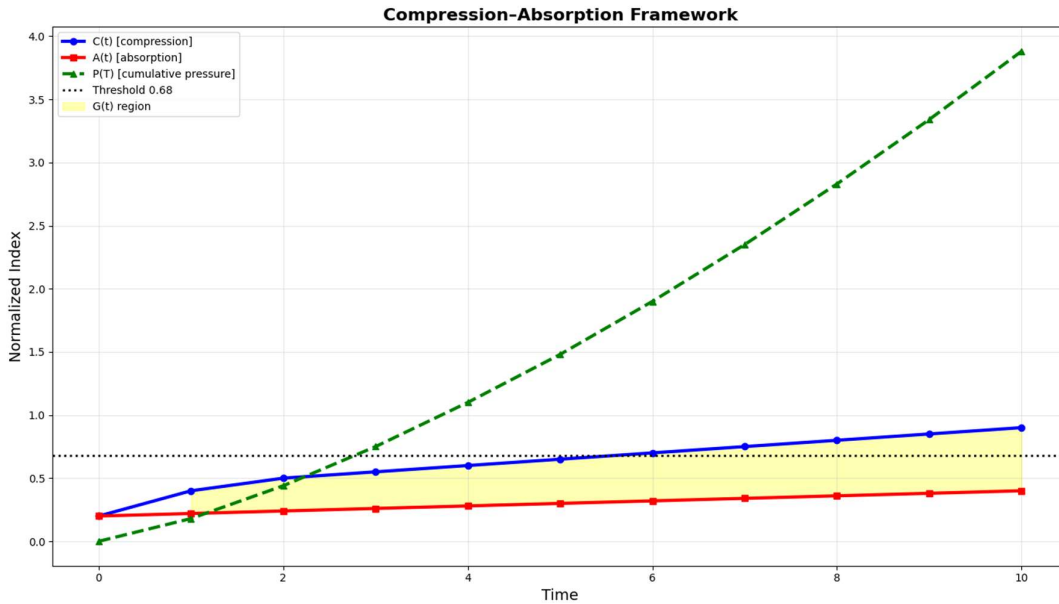


Figure A1. Schematic representation of the framework's core dynamics. The diagram illustrates three phases: (a) stable equilibrium where $A(t)$ matches or exceeds $C(t)$; (b) an opening governance gap where $C(t)$ begins to exceed $A(t)$, generating positive $G(t)$; and (c) lock-in risk accumulation as $P(T)$ rises with sustained positive $G(t)$. The vertical axis represents composite index values (0–1 normalized scale); the horizontal axis represents time. The preemption window is the interval between $G(t)$ first becoming positive and $P(T)$ crossing the estimated entrenchment threshold. Values are illustrative. $C(t)$: compression velocity; $A(t)$: absorptive capacity; $G(t)$: governance gap; $P(T)$: cumulative lock-in pressure.

Note: Values are illustrative. Appendix A4 provides a worked numerical example with specific index values.

A2. Measuring Institutional Absorption Capacity $A(t)$

Absorptive capacity is dynamic and endogenous. General-purpose technology literature emphasizes that institutional complements determine whether productivity gains translate into broad-based prosperity (Brynjolfsson et al., 2021). Accordingly, $A(t)$ must incorporate adaptive institutional variables.

Proposed Components of $A(t)$:

1. **Retraining Velocity**
Percentage of displaced workers successfully reallocated within a defined time horizon.
2. **Sectoral Expansion Rate**
Net job creation in emerging industries relative to displaced positions.
3. **Labor Mobility Elasticity**
Cross-sector occupational transition rates.
4. **Fiscal Buffer Strength**
Capacity of public finance systems to sustain transition support.
5. **Administrative Adaptation Speed**
Time required to enact policy adjustments or implement workforce reforms.
6. **Entrepreneurial Formation Rate**
Growth of new firm creation in AI-complementary sectors.

Absorption is not synonymous with job creation. It reflects systemic capacity to integrate technological acceleration without persistent participation decline.

A3. The Governance Gap $G(t)$

Recall the stability condition:

$$C(t) \leq A(t)$$

Define the governance gap:

$$G(t) = \max \{0, C(t) - A(t)\}$$

When $G(t) > 0$, compression exceeds absorption.

Short-term deviations may be manageable. However, sustained divergence generates cumulative compression pressure:

$$P(T) = \int_0^T G(t) dt$$

Increasing returns theory suggests that cumulative structural shifts become progressively resistant to reversal (Arthur, 1989; Pierson, 2000). As $P(T)$ increases, correction cost likely rises non-linearly, though precise functional relationships require empirical calibration. The candidate functional form and its theoretical grounding are developed in the note below.

A3 Note: Candidate Functional Form for the P(T)–Lock-In Relationship

The relationship between cumulative pressure $P(T)$ and lock-in probability $\lambda(T)$ cannot be formally derived from existing literature. Arthur (1989) and Pierson (2000) establish increasing returns qualitatively; they do not yield a specific probability distribution. This note proposes a candidate functional form chosen for tractability—an S-curve with the qualitative properties the theory requires—not as a derivation from those literatures. It operationalizes Falsification Condition B.1 (Section 7.6) and should be treated as an illustrative conjecture pending empirical calibration. It operationalizes Falsification Condition B.1 (Section 7.6): if reversibility proves cheap and common—contradicting the convex shape implied below—the framework's lock-in predictions would be disconfirmed. The form is not presented as empirically validated; rather, it specifies the form that future calibration research should attempt to estimate.

Proposed form: $\lambda(T) = 1 - \exp(-\alpha \cdot P(T)^\beta)$, where $\alpha > 0$ scales the rate at which accumulated pressure converts to lock-in probability, and $\beta > 1$ imposes convexity — meaning the marginal contribution of each additional unit of $P(T)$ to lock-in probability increases as $P(T)$ grows. This is a standard increasing, bounded, convex-then-concave form (an S-curve with inflection at $P^* = (\beta - 1)/(\alpha\beta)^{1/\beta}$) that captures three substantive features the theory requires. First, at low $P(T)$, lock-in probability rises slowly: early governance gaps are recoverable at modest cost, consistent with the correction cost window analysis in §3.3. Second, above the inflection point P^* , the curve steepens sharply: accumulated structural reorganization (capital deepening, skill atrophy, political coalition formation around concentrated actors) begins to reinforce itself, consistent with Arthur's (1989) increasing returns mechanism and Pierson's (2000) positive feedback dynamics. Third, as $P(T) \rightarrow \infty$, $\lambda(T) \rightarrow 1$ asymptotically: categorical irreversibility is not asserted — consistent with §3.3's formulation of lock-in as a rising cost trajectory rather than an absolute threshold — but correction becomes prohibitively difficult with probability approaching certainty.

The parameters α and β are jurisdiction-specific and require empirical estimation. Historical automation episodes — particularly the US manufacturing robot adoption wave documented by Acemoglu and Restrepo (2020) — provide the most tractable calibration targets: commuting zones where $P(T)$ accumulated rapidly show persistent non-convergence of wages and employment, consistent with the steep portion of the proposed curve. Calibration would require estimating $P(T)$ from historical $G(t)$ trajectories and regressing these against observable lock-in proxies (employment non-convergence rates, retraining program failure rates, political coalition stability around labor-light production). This constitutes the primary empirical agenda for formalizing the framework's lock-in mechanism beyond its current theoretical conjecture status. Until such calibration is completed, the framework treats the inflection point P^* as an unobserved parameter and relies on the qualitative governance implication that the shape implies: preemption is most cost-effective when $P(T)$ is below P^* , and progressively more costly above it — which is precisely the rationale for acting during the irreversibility window rather than after observable entrenchment.

A4. Indicative Threshold Guidance

Exact numerical thresholds vary by jurisdiction. However, preliminary calibration may consider:

- Sustained compression exceeding absorption by 1% annually as manageable under strong institutional capacity.
- Persistent gaps above 2–3% annually over multi-year periods as indicative of rising structural strain.
- Multi-year gaps exceeding 3–4% as elevated lock-in risk.

These figures are illustrative and require contextual validation. The framework is designed to trigger monitoring and policy review rather than mechanical intervention.

A5. Data Sources and Empirical Calibration

Potential data sources include:

- National labor force surveys
- Occupational classification systems
- Sectoral productivity statistics
- Tax revenue composition data
- Social insurance enrollment trends
- Firm-level automation disclosures

Cross-national comparative studies may further calibrate compression–absorption ratios.

Future empirical research should model:

- Functional relationships between $P(T)$ and institutional rigidity
 - Lag structures between task substitution and wage polarization
 - Political response latency to structural labor compression
-

A6. Limitations of Quantification

The Compression–Absorption Framework is diagnostic rather than predictive. Measurement challenges include:

- Attribution difficulty (automation vs. globalization vs. demographic change)
- Time lag between deployment and labor impact
- Informal sector dynamics
- Cross-border technological diffusion

Quantification should therefore be interpreted probabilistically, not deterministically.

The purpose of operationalization is not to forecast collapse but to identify emerging velocity asymmetries before irreversibility thresholds are crossed.

A4. Worked Illustrative Example: Index Construction and Gap Calculation

The following illustrative example demonstrates how C(t), A(t), and G(t) can be constructed from component indicators, how weights are assigned, and how the governance gap is calculated. All values are hypothetical and chosen to demonstrate the logic of the framework, not to assert empirical claims about any specific jurisdiction or period.

Step 1: Construct C(t) — Compression Index

C(t) is a normalized composite index scaled 0–1, constructed from three component indicators. Each component is normalized against a sector-specific baseline and assigned an equal weight of 1/3 in the base specification. Normalization uses a min-max function: $\text{score} = (\text{observed value} - \text{sector minimum}) \div (\text{sector maximum} - \text{sector minimum})$, where the bounds are derived from the historical range of that indicator within the relevant sector or jurisdiction. This converts heterogeneous raw units (percentages, ratios, index values) onto a common 0–1 scale before aggregation. Weights may be adjusted through sensitivity analysis in applied research; equal weighting is used here for transparency and replicability.

C(t) Component	Observed Value	Normalized Score	Weight
Task substitution rate (% routine tasks automated in sector per annum)	8% p.a.	0.80	1/3
Decision-layer entry rate decline (% YoY reduction in new hires for mid-tier analytical roles)	-12% YoY	0.72	1/3
Wage polarization index (ratio of top-decile to median wage growth, sector-specific)	1.8× baseline	0.60	1/3
C(t) composite	—	0.71	(0.80+0.72+0.60)/3

Step 2: Construct A(t) — Absorption Index

A(t) is constructed analogously on the same 0–1 scale, using three absorption indicators. Higher values indicate stronger institutional absorptive capacity. Units are again normalized, enabling direct comparison with C(t) on a common scale.

A(t) Component	Observed Value	Normalized Score	Weight
Retraining velocity (% displaced workers successfully reallocated within 18 months)	35%	0.35	1/3
Fiscal absorption capacity (retraining expenditure as % of GDP relative to OECD median)	0.4× OECD median	0.40	1/3
Labor mobility index (cross-sector transition rate within 24 months of displacement)	28%	0.28	1/3
A(t) composite	—	0.34	(0.35+0.40+0.28)/3

Step 3: Calculate G(t) — Governance Gap

$G(t) = C(t) - A(t) = 0.71 - 0.34 = 0.37$. Both indices are scaled 0–1; G(t) is therefore expressed in the same unit-free normalized scale. A positive G(t) of 0.37 indicates that compression substantially exceeds absorptive capacity — the $C(t) \leq A(t)$ stability condition is violated. Under the framework’s threshold logic, this signals elevated structural lock-in risk and would trigger the preemptive governance response outlined in Section 8. A G(t) at or below zero would indicate that absorption is keeping pace with compression, and no preemptive intervention is warranted at that measurement point.

This example resolves the index construction problem in three respects. First, both C(t) and A(t) operate on a common normalized 0–1 scale, making G(t) directly interpretable without unit conversion. Second, equal weighting is used as the default specification, with sensitivity analysis recommended when domain-specific theory justifies differential weighting. Third, the example makes explicit that retraining velocity is one among three A(t) components — not the sole absorption measure — which prevents single-indicator gaming of the threshold condition. Full operationalization for a specific jurisdiction would require empirical calibration of the normalization bounds (the sector-specific baselines against which observed values are scored), a task addressed in the research agenda at §6.2.

Appendix B: High-Autonomy AI Classification Framework

Institutional preemption requires clear identification of systems capable of generating cross-sector labor compression and authority redistribution. Classification is not prohibition. It is a risk-tiering mechanism designed to determine when deployment velocity warrants macro-institutional review.

This appendix proposes a structured classification framework grounded in institutional risk theory and increasing returns dynamics (Arthur, 1989; Pierson, 2000).

B1. Defining High-Autonomy Systems

A system qualifies as high-autonomy when it satisfies one or more of the following structural characteristics:

1. Decision-Layer Substitution

The system performs multi-step analytical, managerial, or adjudicative functions without real-time human authorization.

Examples include:

- Automated credit approval systems
- Autonomous regulatory enforcement algorithms
- AI-driven hiring and promotion screening
- Healthcare triage systems

Task-based labor research indicates that substitution at the decision layer has broader spillover effects than routine-task automation (Autor, 2013; Acemoglu & Restrepo, 2018).

2. Cross-Sector Scalability

The system can be deployed across multiple industries with minimal adaptation cost, characteristic of general-purpose technologies (Brynjolfsson et al., 2021).

Cross-sector scalability increases compression velocity and amplifies cumulative structural effects.

3. Organizational Authority Integration

The system materially influences or replaces hierarchical coordination functions, such as:

- Strategic planning
- Resource allocation
- Risk modeling
- Compliance adjudication

Authority integration differentiates high-autonomy systems from assistive tools.

4. Autonomous Iterative Optimization

The system updates and refines decision logic dynamically without continuous human recalibration, thereby increasing independence from direct oversight.

Such systems may accelerate path dependency effects by embedding algorithmic standards into institutional processes (Arthur, 1989).

5. Systemic Penetration Threshold

Deployment scale exceeds predefined economic or population thresholds, such as:

- Affecting more than X% of national workforce
- Governing decisions for more than Y% of regulated entities
- Processing more than Z% of sectoral transactions

Thresholds should be calibrated nationally.

B2. Risk-Tier Categories

Classification may adopt a tiered model:

Tier I – Assistive Systems

- Human-in-the-loop remains primary decision authority.
- Limited task substitution.
- Minimal cross-sector scalability.

Regulatory response: Standard compliance.

Tier II – Augmentative Systems

- Partial decision-layer substitution.
- Moderate scalability.
- Human override exists but is not primary workflow driver.

Regulatory response: Monitoring and reporting requirements.

Tier III – High-Autonomy Systems

- Multi-step decision-layer automation.
- Cross-sector deployment capacity.
- Organizational authority integration.
- Potential macro-level labor compression impact.

Regulatory response: Threshold-triggered review under institutional preemption doctrine.

B3. Governance Trigger Mechanism

Classification alone does not impose restriction. It activates structured review under the Compression–Absorption Framework.

Trigger conditions may include:

- Projected compression exceeding defined percentage of sectoral employment.
- Demonstrated cumulative task substitution trends.
- Authority displacement in high-impact civic domains.

The objective is velocity assessment, not innovation suppression.

B4. Authority-Sensitive Domains

Regardless of tier, systems deployed in the following domains require enhanced scrutiny due to legitimacy implications (Lipset, 1959; Fukuyama, 2014):

- Legal adjudication
- Public benefits eligibility
- Financial access determination
- Immigration and citizenship status
- Electoral administration
- Regulatory enforcement

These domains directly implicate sovereignty and institutional durability.

B5. Anti-Arbitrage Safeguards

Classification must apply based on functional impact rather than formal corporate structure. Firms should not evade Tier III designation through fragmentation of deployment across subsidiaries or jurisdictions.

Path dependency literature underscores the importance of early categorization before entrenchment effects accumulate (Pierson, 2000).

B6. International Coordination Considerations

In open economies, classification standards benefit from regional harmonization to reduce regulatory arbitrage (Rodrik, 2011). Coordination need not eliminate national discretion but should align core definitions of high-autonomy systems.

B7. Limitations

Classification requires periodic reassessment. Technological capability evolves, and systems may migrate across tiers over time. Moreover, high-autonomy designation does not imply inherent harm. It signals structural impact potential under velocity asymmetry conditions.

The framework is diagnostic, not prohibitory.

Appendix C: Malaysia and ASEAN Calibration Model

Applying the Compression–Absorption Framework in Open Mid-Sized Economies

Open mid-sized economies operate under structural constraints distinct from large technological powers. They do not set global deployment standards, yet they are highly exposed to cross-border capital mobility, supply chain integration, and imported technological architectures. Malaysia and ASEAN member states exemplify this configuration.

This appendix proposes a calibration model for applying the Compression–Absorption Framework in such environments.

C1. Structural Characteristics of Open Mid-Sized Economies

Malaysia and many ASEAN economies share several structural features:

1. **Export-Dependent Growth Models**
Manufacturing and trade-intensive sectors contribute significantly to GDP.
2. **Integration into Global Value Chains**
Technological adoption may be externally driven rather than domestically initiated.
3. **Young or Transitioning Labor Forces**
Entry-level absorption capacity has heightened political sensitivity.
4. **Limited Unilateral Regulatory Leverage**
Major AI platforms and system providers may be foreign-based.
5. **Fiscal Constraints**
Limited capacity for prolonged income stabilization compared to advanced welfare states.

These characteristics amplify exposure to acceleration asymmetry (Rodrik, 2011).

C2. Sectoral Calibration Under the Two-Speed Governance Model

The Two-Speed Governance Model should be calibrated regionally as follows:

Competitive Acceleration Lane

Faster deployment may be tolerated in:

- Semiconductor manufacturing
- Industrial automation
- Logistics optimization
- Export-oriented services
- Cybersecurity infrastructure

In these sectors, productivity gains and global competitiveness may offset labor compression effects, provided absorption capacity expands concurrently.

Stability Preservation Lane

Stricter thresholds should apply in sectors central to civic legitimacy and social cohesion:

- Public administration
- Welfare eligibility systems
- Credit and financial access determination
- Healthcare triage
- Legal adjudication
- Education certification systems

These sectors implicate sovereignty, authority chains, and political trust (Fukuyama, 2014; Lipset, 1959).

Velocity modulation here is protective of institutional equilibrium.

C3. Indicative Calibration Thresholds (Malaysia Context)

Illustrative parameters for national review triggers may include:

- Sectoral automation exceeding 15–20% of decision-layer roles within a five-year horizon.
- Entry-level employment contraction exceeding 10% in managerial or analytical occupations.
- Sustained divergence between productivity growth and median wage growth beyond historical norms.
- Labor share decline exceeding established historical volatility bands.

These thresholds are indicative and require empirical refinement through national statistical agencies.

C4. ASEAN Coordination Mechanism

Regional coordination reduces destabilizing regulatory arbitrage. ASEAN-level dialogue may focus on:

- Shared classification criteria for high-autonomy systems.
- Common reporting standards for compression indicators.
- Information exchange on cross-border deployment trends.
- Soft harmonization of authority-sensitive domain protections.

Such coordination is consistent with regional integration models that preserve sovereignty while mitigating competitive races (Rodrik, 2011).

C5. Youth Employment and Intergenerational Risk

In emerging and middle-income economies, entry-level labor absorption is closely tied to upward mobility expectations. Prolonged compression in decision-layer occupations may disproportionately affect educated youth cohorts.

Political legitimacy theory suggests that when upward mobility expectations weaken, trust erosion may follow (Lipset, 1959). Early monitoring of entry bottlenecks is therefore critical.

C6. Fiscal Adaptation Strategy

Malaysia and ASEAN states may lack capacity for indefinite universal income stabilization. Therefore, absorption expansion should prioritize:

- Targeted retraining subsidies
- SME digital augmentation grants
- Entrepreneurship ecosystem support
- Regional mobility facilitation
- Sectoral reallocation incentives

These measures increase $A(t)$ without permanently expanding fiscal burdens.

C7. Strategic Positioning

Institutional preemption in open economies should not signal technological hesitancy. Rather, it positions states as:

- Competitively agile
- Institutionally resilient
- Governance-forward

Stability under rapid transformation can itself become a competitive advantage.

C8. Limitations of Regional Application

Calibration parameters must reflect:

- Informal sector dynamics
- Cross-border labor mobility
- External capital influence
- Political system variation within ASEAN

The model is adaptive and must evolve alongside technological deployment patterns.

Appendix D: Human Agency Safeguard Protocol

Preserving Accountable Authority in High-Impact Domains

The doctrine of Institutional Preemption identifies the preservation of accountable human authority as a structural requirement of modern governance. This appendix operationalizes that requirement through a Human Agency Safeguard Protocol (HASP) applicable to high-autonomy systems deployed in authority-sensitive domains.

The objective is not to prohibit automation. It is to preserve identifiable responsibility, traceable authority chains, and corrective institutional capacity (North, 1990; Fukuyama, 2014).

D1. Rationale

Modern legal and administrative systems depend on:

1. Identifiable decision-makers
2. Clear liability attachment
3. Procedural review pathways
4. Democratic accountability mechanisms

Economic participation and distributed authority have long been linked to political legitimacy (Lipset, 1959). When decision-layer authority becomes computationally centralized without meaningful human override, institutional coherence may weaken.

The Human Agency Safeguard Protocol seeks to prevent such structural authority migration.

D2. Authority-Sensitive Domains

Enhanced safeguards should apply where decisions directly affect:

- Legal status (adjudication, sentencing)
- Financial access (credit approval, loan denial)
- Public benefits eligibility
- Healthcare triage and allocation
- Immigration and citizenship status
- Regulatory enforcement actions
- Electoral administration

These domains implicate sovereignty, rights, and civic trust.

D3. Core Safeguard Requirements

1. Identifiable Human Decision Authority

A named, accountable human official must retain final decision authority.

This includes:

- Authority to override algorithmic output
- Obligation to document reasoning
- Personal liability within institutional structure

Symbolic human presence without real override capacity does not satisfy the protocol.

2. Auditability and Traceability

Systems must maintain:

- Decision logs
- Version histories
- Model update records
- Input-output traceability

Auditability preserves review capacity and legal defensibility (North, 1990).

3. Procedural Appeal Pathways

Affected individuals must have access to:

- Clear explanation mechanisms
- Administrative review processes
- Judicial recourse where applicable

Without appeal mechanisms, authority shifts from accountable governance to opaque systems.

4. Proportional Autonomy Constraints

Autonomy level should be proportional to domain sensitivity.

For example:

- Assistive tools may operate with minimal restriction in low-impact domains.
- Fully autonomous systems should not determine legal status or fundamental rights without substantive human adjudication.

This aligns with tiered classification principles in Appendix B.

5. Concentration Safeguards

Where a single computational system governs decisions affecting large portions of the population, systemic risk increases.

Safeguards may include:

- Institutional redundancy
- Independent oversight bodies
- Periodic authority review

Increasing returns theory suggests concentration amplifies entrenchment risk (Arthur, 1989; Pierson, 2000).

D4. Preventing Authority Hollowing

A key risk is authority hollowing: formal human responsibility persists while substantive discretion shifts to algorithmic outputs.

To prevent hollowing:

- Human decision-makers must receive independent training.
- Institutional culture must reinforce discretionary responsibility.
- Performance metrics must not penalize human override of algorithmic recommendations.

Human agency must be operational, not ceremonial.

D5. Sovereignty Implications

Sovereignty in modern states depends on:

- Traceable chains of command
- Distributed economic participation
- Institutional capacity for corrective adjustment (Rodrik, 2011; Fukuyama, 2014)

If decision-layer authority centralizes into computational systems controlled by limited actors—particularly transnational providers—sovereignty may effectively migrate beyond traditional institutional boundaries.

Human agency safeguards thus function as constitutional stabilizers, preserving corrective optionality within governance systems.

D6. Relationship to Innovation

The Human Agency Safeguard Protocol does not restrict innovation in:

- Industrial optimization
- Scientific research
- Logistics efficiency
- Creative augmentation

It applies only where authority intersects with rights, legitimacy, and institutional durability.

The goal is not technological restraint, but institutional continuity.

D7. Limitations

Human agency safeguards cannot eliminate all systemic risk. Overly rigid requirements may:

- Slow administrative efficiency
- Increase compliance costs
- Generate procedural burdens

Therefore, safeguards should be calibrated proportionally and reviewed periodically.

Preserving agency requires balance, not absolutism.

Appendix E: Empirical Research Agenda

Advancing the Study of Labor Compression and Institutional Absorption

The Compression–Absorption Framework articulated in this paper is structural and diagnostic. Its maturation requires empirical refinement across jurisdictions, sectors, and technological domains. This appendix outlines a research agenda to operationalize and test the framework.

E1. Measuring Compression Velocity

Future research should quantify labor compression velocity across sectors, distinguishing between:

- Task substitution within occupations (Autor, 2013; Acemoglu & Restrepo, 2018)
- Occupational elimination
- Wage polarization effects
- Entry-level bottleneck formation

Key research questions:

1. What is the lag between high-autonomy deployment and observable wage compression?
2. Does decision-layer automation produce faster displacement effects than routine-task automation?
3. Are cross-sector compression patterns synchronized under AI diffusion?

Longitudinal datasets linking firm-level automation adoption to labor market outcomes are critical.

E2. Modeling Absorptive Capacity Dynamics

Absorptive capacity $A(t)$ is endogenous and may expand through complementary institutional investment (Brynjolfsson et al., 2021).

Future research should examine:

- Elasticity of retraining systems relative to displacement magnitude.
- Cross-sector mobility constraints.
- Fiscal sustainability of transition support mechanisms.
- Institutional reform latency under technological shock.

Comparative cross-country studies could test whether higher baseline institutional capacity moderates compression effects.

E3. Estimating the Governance Gap

Operationalizing the governance gap $G(t) = C(t) - A(t)$ requires empirical calibration.

Research priorities include:

1. Developing composite indices for C and A.
2. Identifying threshold levels associated with persistent labor share decline.
3. Estimating cumulative compression pressure $P(T)$ in historical analogues (e.g., industrial automation waves).

Increasing returns theory suggests non-linear entrenchment effects (Arthur, 1989; Pierson, 2000), but empirical confirmation remains limited.

E4. Decision-Layer Automation and Authority Distribution

A distinct empirical question concerns authority redistribution.

Research should examine:

- Degree of human override exercised in AI-mediated decision systems.
- Impact of algorithmic decision-making on institutional accountability.
- Correlation between authority centralization and public trust indicators.

Institutional durability literature suggests that legitimacy depends on accountable authority structures (Lipset, 1959; Fukuyama, 2014). Empirical validation in AI-mediated contexts is necessary.

E5. Cross-Border Deployment and Sovereignty Effects

Open economies face cross-border technological diffusion.

Research questions:

- Does foreign ownership of high-autonomy systems affect national policy autonomy?
- How does regulatory arbitrage influence compression velocity?
- Does regional coordination moderate structural entrenchment?

These questions intersect with globalization constraints identified by Rodrik (2011).

E6. Political Response Latency

Political systems frequently respond to visible crisis rather than anticipatory indicators (Pierson, 2000).

Future work should explore:

- Median lag between measurable compression signals and policy intervention.
- Determinants of anticipatory governance adoption.
- Comparative case studies of early versus delayed response.

Understanding response latency is essential for operationalizing institutional preemption.

E7. Intergenerational Effects

Entry-level absorption dynamics warrant particular study in emerging economies.

Research should assess:

- Youth unemployment sensitivity to decision-layer compression.
- Long-term income scarring from early displacement.
- Political and social trust impacts among affected cohorts.

Such analysis is especially relevant for regions with youthful labor demographics.

E8. Simulation and Scenario Modeling

Agent-based modeling and macroeconomic simulation can test:

- Non-linear lock-in dynamics.
- Interaction between compression and absorption variables.
- Fiscal sustainability under varying compression trajectories.

These models would refine the functional relationship between cumulative pressure $P(T)$ and institutional rigidity.

E9. Normative Evaluation

Finally, research should evaluate:

- Whether preserved human authority measurably improves institutional trust.
- Trade-offs between efficiency gains and agency safeguards.
- Optimal calibration of velocity modulation thresholds.

Normative evaluation must remain grounded in institutional function rather than philosophical abstraction.

E10. Conclusion of Research Agenda

The Compression–Absorption Framework establishes a structural hypothesis:

Sustained acceleration asymmetry increases entrenchment probability.

Testing this hypothesis requires interdisciplinary collaboration across labor economics, institutional theory, political science, public finance, and AI governance.

The objective is not to validate alarmist predictions but to determine empirically whether velocity management is necessary to preserve corrective optionality in the age of labor-compressing AI.

Appendix F: Constitutional and Sovereignty Implications of High-Autonomy AI

Authority, Legitimacy, and the Migration of Decision Power

The Compression–Absorption Framework identifies labor compression as an economic variable. However, when compression occurs within decision-layer domains—credit allocation, regulatory enforcement, welfare eligibility, risk assessment, adjudication—its implications extend beyond labor markets into constitutional structure. This appendix examines how labor-compressing AI may affect sovereignty, authority distribution, and institutional reversibility.

F1. Sovereignty as Institutional Authority

In modern states, sovereignty is not merely territorial control. It consists of:

1. Identifiable chains of authority
2. Distributed economic participation
3. Capacity for corrective institutional adjustment
4. Democratic or administrative accountability mechanisms

Institutional theory emphasizes that authority must remain traceable to maintain legitimacy and stability (North, 1990; Fukuyama, 2014). Political legitimacy has also been historically linked to broad economic participation and distributive integration within the state (Lipset, 1959).

When decision-layer functions are increasingly mediated by computational systems operating at scale, sovereignty may shift functionally—even where formal legal authority remains unchanged. The question is not whether law disappears, but whether operational authority becomes structurally concentrated beyond ordinary channels of contestation.

F2. Authority Migration and Computational Centralization

A growing body of scholarship documents how algorithmic systems structure decisions in domains such as credit scoring, predictive policing, welfare administration, hiring, insurance underwriting, and sentencing risk assessment (Pasquale, 2015; Crawford, 2021). Even where human officials retain formal responsibility, algorithmic outputs may meaningfully shape, constrain, or pre-structure discretion.

Bovens and Zouridis (2002) describe the evolution from “street-level bureaucracy” to “system-level bureaucracy,” in which discretion migrates from frontline officials into embedded rule systems. Labor-compressing AI extends this trajectory by embedding adaptive, data-driven decision logic into scalable architectures.

This dynamic may be described as **authority migration**: operational decision power shifts from distributed human agents to centralized computational infrastructures. Increasing returns theory suggests that once such systems become embedded, reversal may become progressively more costly as complementary investments, organizational routines, and skill profiles adapt around them (Arthur, 1989; Pierson, 2000).

Authority migration does not require malicious intent. It may emerge incrementally through efficiency optimization, risk management, and standardization pressures.

F3. Constitutional Hardening Without Formal Amendment

Labor-compressing AI may reshape governance practices without formal constitutional amendment (Pasquale, 2015). Legal texts can remain unchanged while:

- Decision authority becomes algorithmically mediated
- Review capacity becomes technically constrained
- Administrative discretion narrows
- System designers acquire structural influence

Bovens and Zouridis (2002) characterize such developments as shifts toward system-level administration. Crawford (2021) further highlights how computational infrastructures embed material and political power within technological architectures.

In this sense, production architecture may acquire quasi-constitutional effects. Governance practices may harden functionally—through code, data pipelines, and optimization routines—without explicit legal reform. The risk is not necessarily authoritarian rupture, but gradual institutional inflexibility.

F4. Corrective Optionality as a Constitutional Principle

The main text introduced **corrective optionality**: the institutional capacity to recalibrate authority and participation structures before entrenchment.

At the constitutional level, corrective optionality implies:

- Preservation of meaningful human override authority
- Maintenance of substantive appeal pathways
- Transparency sufficient for contestation
- Prevention of excessive concentration of decision-layer control

Citron and Pasquale (2014) argue that technological due process requires more than symbolic review; it requires genuine capacity to understand, challenge, and reverse automated determinations. Path dependency literature further suggests that early structural commitments may constrain future policy space (Pierson, 2000).

Where algorithmic systems become deeply embedded in authority-sensitive domains, delayed oversight may narrow constitutional maneuverability.

F5. Transnational Platform Sovereignty

In open economies, many high-autonomy systems are developed and maintained by transnational corporations. When domestic decision-layer functions depend on externally controlled computational infrastructures, sovereignty may become partially externalized.

Crawford (2021) documents concentration of data resources, compute capacity, and algorithmic expertise among a small number of global actors. Rodrik (2011) demonstrates how globalization can constrain domestic policy autonomy under deep integration conditions.

Labor-compressing AI may intensify these constraints if governance-critical functions—credit allocation, benefits administration, regulatory analytics—rely on systems governed beyond domestic institutional control.

Institutional preemption therefore intersects with digital sovereignty and policy autonomy concerns.

F6. Human Agency as Constitutional Stabilizer

The Human Agency Safeguard Protocol (Appendix D) functions not as resistance to automation, but as structural stabilization.

Preserving:

- Named decision authority
- Substantive override capacity
- Transparent review mechanisms

ensures that authority remains contestable and institutionally anchored.

Human agency in this framework is not a metaphysical claim about human superiority. It is a constitutional design principle. It maintains traceability, liability attachment, and the possibility of institutional correction (Fukuyama, 2014).

Without identifiable authority chains, governance coherence weakens—even if efficiency increases.

F7. Gradual Drift Versus Crisis

Historical institutional change frequently proceeds incrementally rather than through rupture (North, 1990; Pierson, 2000). Structural drift may advance through routine optimization rather than overt political decision.

Labor-compressing AI may accelerate such drift by embedding decision logic within scalable computational architectures (Pasquale, 2015; Crawford, 2021). Unlike infrastructure-heavy technological shifts that required physical reorganization, algorithmic systems can be deployed, updated, and scaled with relatively lower institutional friction.

The constitutional question is therefore temporal:

Can institutional adaptation keep pace with authority migration?

F8. Limits and Counterarguments

Several counterarguments warrant consideration:

1. Formal accountability may persist even where operational mediation increases.
2. Post-wage or hybrid economic systems may redefine participation norms.
3. Decentralized or open AI systems may distribute rather than concentrate authority.

The literature documents patterns of algorithmic opacity and concentration (Eubanks, 2018; Crawford, 2021), but technological trajectories remain contested.

The framework identifies structural risk conditions, not deterministic outcomes.

F9. Concluding Observation

AI governance debates often center on safety, bias, or innovation policy. This appendix suggests that labor-compressing AI may also implicate constitutional structure.

When decision-layer authority migrates into scalable computational systems, sovereignty dynamics may shift gradually and often invisibly.

Institutional preemption does not seek to halt technological transformation. It seeks to ensure that authority, accountability, and corrective optionality remain within democratic and administrative reach.

The preservation of sovereignty in the age of labor-compressing AI is not a question of resisting machines. It is a question of governing velocity before structural entrenchment narrows the space for correction.
