Too Fast to Adjust

Adoption Speed and the Permanent Cost of AI Transitions

March 2026

Eduardo Levy Yeyati
Universidad Torcuato Di Tella - IDB

Abstract

We study how the speed of AI adoption affects labor market outcomes during technological transitions. In a dynamic model where displaced routine workers enter a retraining pipeline with finite capacity, faster adoption compresses the displacement window without reducing total displacement, overwhelming the pipeline and generating permanent labor force exit through worker discouragement. The central result is that, even when two economies share the same long-run automation level, adoption speed alone determines transition welfare: faster adoption produces a larger discouraged stock, a steeper and more persistent decline in labor force participation, and a sustained compression of the labor share throughout the transition window. Non-routine employment and wages exhibit a crossing pattern — initially higher under fast adoption, then lower — so that faster adoption can simultaneously raise long-run wages for survivors while permanently reducing participation. Social welfare is strictly concave in adoption speed and maximized at an interior optimum below the market rate, because firms do not internalize the congestion externality they impose on the retraining queue, the irreversibility of permanent exit, or the wage depression borne by non-routine incumbents. The socially optimal speed and retraining capacity are complements: stronger institutions raise the optimal adoption speed.

Executive Summary

The debate over artificial intelligence and jobs has focused almost entirely on the wrong variable. Policy, research, and public commentary have concentrated on *how much* gets automated — which tasks, which industries, which workers. This paper argues that *how fast* matters at least as much, and potentially more. Two economies that end up with the same level of automation can experience radically different labor market outcomes during the transition. The one that gets there faster is not necessarily better off — it may be permanently worse off, even if its long-run productivity is identical.

When automation displaces routine workers, those workers enter a retraining pipeline with limited capacity. Fast adoption compresses this displacement into a shorter window. The same total number of displaced workers arrives in a fraction of the time, overwhelming the retraining system. Workers who rationally assess their prospects — long queues, depressed wages, uncertain re-employment — exit the labor force permanently rather than wait. This exit is absorbing. The welfare cost of fast adoption is therefore not captured by where labor force participation ends up, but by how steeply it falls and how long it stays depressed, and by the sustained compression of the labor share throughout the transition window.

The paper establishes four results. Peak unemployment is strictly increasing in adoption speed: the same total displacement arrives faster, so the peak rises even though aggregate displacement is unchanged. Permanent labor force exit is larger under fast adoption even when terminal participation levels are similar across paths — the discounted labor share is unambiguously lower throughout the transition. Non-routine employment and wages follow a crossing pattern: initially higher under fast adoption as the compressed displacement delivers retrained workers sooner, then lower as the surge of retrainees depresses wages and the larger discouraged stock depletes the employed workforce. Faster adoption can simultaneously raise long-run wages for survivors while reducing overall participation — an outcome that wage-level statistics alone would mischaracterize as beneficial. Finally, social welfare is maximized at an interior adoption speed strictly below the market rate, because individual firms do not internalize the congestion they impose on other workers' retraining queues, the irreversibility of permanent exit, or the wage depression imposed on non-routine incumbents.

1 Introduction

For most of the past decade, AI adoption has been organizationally incremental: tools layered onto existing firms, accelerating tasks without restructuring the production processes around which those tasks are organized. Recent developments suggest this is changing.

In early 2026, Project Prometheus—a start-up led by Jeff Bezos and former Google executive Vikram Bajaj, valued at $30bn after a $6.2bn raise—announced plans to raise tens of billions more through a holding company targeting industrial firms disrupted by AI.[1] Prometheus is not selling AI tools to manufacturers. It is building AI systems capable of mapping physical processes, understanding engineering design, and modeling the manufacturing of complex objects—jet engines, semiconductors—with the explicit aim of reorganizing those industries around what it builds. The capital scale and the sovereign wealth fund backing illustrate what a high-κ adoption episode looks like: organizational restructuring at scale, not incremental augmentation of existing workflows.

The historical parallel is precise. The steam engine did not transform manufacturing when it was invented; it did so when firms restructured their factory layouts around it. In 1913–14, Henry Ford reorganized Highland Park around the moving assembly line. Within eighteen months, output per worker tripled and annual worker turnover reached 370 percent. The technology was unambiguously productivity-enhancing; the transition was not. Ford's response—the $5 day, more than doubling wages overnight—was not philanthropy. It was the market's belated correction for absorbing displacement faster than the Detroit labor market could reallocate workers on its own.

Automation episodes like these may follow the same sequence as 1913: incremental adoption gives way to structural reorganization, and the labor market cost depends not on how much gets automated—the long-run endpoint—but on *how fast*. This paper formalizes that argument. Holding fixed the long-run automation envelope, transition welfare depends on adoption speed.

We build a dynamic model of labor market transition in which the long-run automation threshold $\bar{\iota}$—the share of tasks eventually automated—is taken as given from the task-based framework of Acemoglu and Restrepo [2018, 2022]. We model the speed κ at which this threshold is reached via logistic diffusion, and study its consequences for unemployment, permanent labor force exit, wages, and

[1] Financial Times, February 27, 2026. The vehicle is in talks with the Abu Dhabi Investment Authority and JPMorgan's supply chain resilience fund.

welfare. The paper does not take a stand on whether automation ultimately raises aggregate productivity or expands the task frontier; it studies a dimension largely abstracted from in steady-state analyses. The central result is: *even if the long-run equilibrium is efficient, the transition need not be.*

The economic mechanism operates through what we call the *shadow price of time in the retraining queue*. Fast adoption compresses the displacement window: the same total mass of displaced workers $\bar{\imath}$ arrives in a shorter period, overwhelming retraining capacity $\rho(Z)$. Workers who rationally update on lengthening queues and depressed non-routine wages exit the labor force permanently rather than waiting for a slot that may not arrive. This exit is absorbing. The welfare cost of fast adoption is therefore not measured by where LFP ends up, but by how steeply it falls and how long it stays depressed—and by the sustained compression of the labor share during the transition window. Figure 1 illustrates: under fast adoption LFP drops sharply within the first decade and the labor share compresses earlier and more steeply, even when terminal levels are similar.

The model features four worker states: employed in routine tasks, employed in non-routine tasks, unemployed and retraining, and permanently discouraged. Peak displacement inflow $\Phi_{\max} = \bar{\imath}\kappa/4$ scales linearly with κ, overwhelming the retraining pipeline. The same long-run automation level $\bar{\imath}$ thus generates vastly different transition paths depending on the κ chosen.

Four propositions characterize the transition:

1. *Unemployment hump* (Proposition 1). Peak unemployment is strictly increasing in κ and $\bar{\imath}$, and strictly decreasing in reallocation capacity $\rho(Z)$. Faster adoption concentrates displacement into a shorter window, raising the peak without changing total displacement $\int \Phi \, dt = \bar{\imath}$.

2. *Permanent exit and LFP path* (Proposition 2). The long-run discouraged stock $D(\infty) \in (0, \bar{\imath})$ is strictly increasing in κ. But the primary welfare object is the *path* of labor force participation $1 - D(t)$: under fast adoption, LFP falls earlier and more steeply, depressing the discounted labor share over the entire transition window even when terminal levels are similar.

3. *Wage depression* (Proposition 3). Faster adoption generates earlier and steeper compression of non-routine wages because retrained workers enter the non-routine sector faster than task demand expands. The discounted cumulative wage depression is strictly increasing in κ.

4. *Interior optimal speed* (Proposition 4). There exists a unique $\kappa^* \in (0, \infty)$

maximizing social welfare. The market speed $\kappa^{\text{mkt}} > \kappa^*$ because firms do not internalize the congestion externality on other workers' retraining queues, the irreversibility of permanent exit, or the wage depression imposed on non-routine incumbents. Calibrations suggest the welfare difference between market-speed and socially optimal adoption is quantitatively significant at plausible parameter values.

Our analysis focuses on the transition phase of AI adoption, when displacement and retraining dynamics dominate and productivity gains are not yet fully realized. In later stages, as AI diffusion raises aggregate output, these forces may attenuate or reverse. The model also has distributional implications: faster adoption can raise wages for employed workers while increasing permanent exit, thereby widening inequality and reducing aggregate labor income.

Note that κ is not a technological constant but the aggregate outcome of decentralized adoption intensity—the speed at which firms collectively choose to integrate, reorganize, and deploy AI systems. Firms choose the intensity and timing of AI integration based on private costs and benefits, internalizing neither the congestion they impose on the retraining pipeline nor the irreversibility of the permanent exit they trigger. κ^{mkt} is therefore a market outcome, not a physical constraint, and it is in principle subject to policy. The comparative statics of κ^* (Corollary 1) identify two distinct instruments. The first is *retraining capacity* $\rho(Z)$—the throughput rate of the labor market's reallocation pipeline, shifted by active labor market programs, retraining subsidies, credential reform, and mobility support. The second is *adoption speed* κ itself, which within limits can be influenced by regulation, AI liability frameworks, sector-specific standards, and the timing of public procurement.

These two levers are *complements, not substitutes*. The central comparative static is $\partial\kappa^*/\partial\rho(Z) > 0$: stronger retraining institutions increase the optimal adoption speed, because a more capable pipeline can absorb higher displacement flows without generating permanent exit. The standard political-economy intuition—protect workers by slowing technology—inverts under our mechanism. The correct policy mix is to invest in retraining capacity *and then let adoption accelerate*, not to choose between them. The planner's problem is not a speed limit but a coordination problem between two complementary investments.

The paper can be read as a dynamic complement to the task-based automation model of Acemoglu and Restrepo [2018, 2022]—compatible with it and, in principle, combinable. That framework characterizes long-run equilibria: which tasks get automated, what new tasks emerge, where wages settle, how the capital–labor

split evolves. These are results about the *destination*. We take the destination as given—$\bar{\imath}$ is the Acemoglu–Restrepo envelope—and ask about the *path*. The distinction matters because comparative statics across steady states are silent on transition dynamics: the welfare cost of getting from one equilibrium to another depends entirely on the speed. Section 7 formalizes this through a two-speed extension in which task creation follows its own logistic at speed κ_T; the Acemoglu–Restrepo benchmark is recovered as $\kappa_T \to \infty$ with an unbounded task frontier.

In a new variation of the well-known Tinbergen race [Tinbergen, 1975, Goldin and Katz, 2008]—technology versus education, with the skill premium as the finish line—the runners here are adoption speed κ and retraining capacity $\rho(Z)$, and the outcome at stake is permanent labor force exit rather than wage inequality. We develop the parallel fully in the discussion.

Section 2 reviews the related literature. Section 3 presents the model. Section 4 defines welfare and the social planner's problem. Section 5 states and proves the main propositions. Section 6 presents testable implications and the research agenda. Section 7 extends to two-speed dynamics. Appendices contain the CES derivation, welfare concavity proof, calibration targets, and existence–uniqueness of equilibrium paths.

2 Related Literature

This paper sits at the intersection of five bodies of work.

Task-based automation. Acemoglu and Restrepo [2018] develop the canonical framework in which automation displaces labor from tasks, and new task creation determines whether employment and labor share recover in the long run. Acemoglu and Restrepo [2022] extend this to account for the full distribution of wage changes across skill groups, matching the rise in US wage inequality since 1980. Our model takes the long-run automation threshold $\bar{\imath}$ and the task creation frontier directly from this framework. The key departure is that we endogenize the transition *speed* κ rather than treating automation as a comparative-static shift. Within the Acemoglu–Restrepo envelope, speed is the sole determinant of transition welfare—a dimension the steady-state framework cannot address. Acemoglu [2002] on directed technical change provides the broader welfare tradition within which our social planner's problem sits: the socially optimal trajectory of technology adoption need not coincide with the market equilibrium, and corrective instruments are welfare-improving.

Adjustment costs and slow labor market reallocation. A central empirical fact motivating our model is that displaced workers do not instantaneously reallocate. Autor and Dorn [2013] document that US local labor markets exposed to Chinese import competition exhibited persistent employment and wage losses lasting over a decade, with substantial permanent exit from the labor force. This finding—that adjustment is slow, costly, and incomplete—is the empirical antecedent to our $D(\infty) > 0$ result. Dix-Carneiro [2014] and Caliendo et al. [2019] provide structural quantification of transition costs in trade liberalization episodes, establishing that dynamic general equilibrium models of labor reallocation generate welfare losses from speed that static models miss entirely. Our model applies this logic to AI adoption: the welfare cost is not in the destination but in the path, and the path cost is increasing in κ.

Labor force participation and permanent exit. Elsby et al. [2013] establish that the participation margin is quantitatively important for understanding labor market fluctuations and is distinct from the unemployment margin. Participation decisions respond to expected re-employment prospects—the same channel formalized here through the discouragement hazard $\mu(q(t), w^N(t))$. Jaimovich and Siu [2012] document that job polarization is concentrated in recessions, when routine employment collapses and does not recover—a pattern consistent with permanent exit being triggered by congestion rather than steady-state displacement. Autor [2019] find that automation is labor-displacing at the industry level, with effects on employment and the labor share that accumulate over time rather than appearing instantaneously. Our model provides a mechanism for this accumulation: the permanent discouraged stock $D(\infty)$ builds along the transition path as a function of κ, not as a feature of the long-run equilibrium.

Technology diffusion and cross-country adoption. Comin and Hobijn [2010] document enormous cross-country variation in both the speed and the eventual penetration of major technologies, with adoption lags of decades for many general-purpose technologies. This empirical regularity justifies treating κ as a parameter that varies across countries and is the basis for our cross-country empirical strategy. Brynjolfsson et al. [2017] argue that AI is a general-purpose technology whose productivity gains are delayed by the organizational complementarities required to exploit it—the same complementarities that determine the timing of the wholesale restructuring we model as a high-κ event. The new task content of employment documented by Autor et al. [2024] provides the empirical basis for the task creation

component κ_T in our two-speed extension.

Optimal policy under technological change. Heathcote et al. [2017] characterize optimal tax progressivity in an economy with skill heterogeneity, establishing that the welfare-maximizing fiscal instrument depends on the elasticity of labor supply across skill groups—a parallel to our finding that the welfare-maximizing κ^* depends on $\rho(Z)$ and σ. The Tinbergen–Goldin–Katz tradition [Tinbergen, 1975, Goldin and Katz, 2008] frames rising inequality as a race between technology and education, with the skill premium as the outcome variable. We propose the analogous race for the current automation episode in the discussion section: adoption speed κ versus retraining capacity $\rho(Z)$, with permanent labor force exit as the finish line. The key structural difference—that the two runners in our race are complements rather than substitutes—is what gives the policy analysis its distinctive character.

3 A Minimal Speed–Capacity Model

Before presenting the full model, we sketch a minimal version that exposes the core mechanism without CES production, existence appendices, or two-speed extensions. The four results below are proved rigorously as Propositions 1–4 in Section 6; this section states them informally so the reader can follow the mechanism before the machinery.

Environment

A unit mass of routine tasks is progressively automated up to a fixed long-run envelope $\bar{\imath} \in (0, 1)$. Let $a(t) \in [0, 1]$ denote the fraction of the envelope already automated. Adoption follows logistic diffusion with speed $\kappa > 0$:

$$\dot{a}(t) = \kappa\, a(t)(1 - a(t)), \qquad a(0) = a_0 \in (0, 1).$$

The instantaneous *displacement flow* is $\Phi(t; \kappa) = \bar{\imath}\, \dot{a}(t)$, so $\Phi_{\max} = \bar{\imath}\kappa/4$ and $\int_0^\infty \Phi\, dt = \bar{\imath}(1 - a_0)$: higher κ compresses the same total displaced mass into a shorter window.

Labor market

Displaced workers enter unemployment $U(t)$ and face two competing hazards: retraining success at rate $\rho > 0$ (institutional capacity, taken as exogenous here) and permanent exit at endogenous rate $\mu(t) \geq 0$ (discouragement). Let $L_N(t)$ denote successfully retrained workers and $D(t)$ the permanently inactive stock:

$$\dot{U}(t) = \Phi(t;\kappa) - (\rho + \mu(t))\, U(t),$$
$$\dot{L}_N(t) = \rho\, U(t), \qquad \dot{D}(t) = \mu(t)\, U(t).$$

Non-routine wages clear a reduced-form labor demand:

$$w_N(t) = w_{N0} \left(\frac{L_{N0}}{L_N(t)} \right)^{1/\sigma}, \qquad L_{N0} = 1 - \bar{\imath}, \quad \sigma > 0,$$

where L_{N0} is the initial non-routine workforce. As retrained workers enter L_N, wages fall; as permanent exit depletes L_N in the long run, wages recover. Discouragement responds to both congestion and wages:

$$\mu(t) = \max(0,\ \mu_0 + \mu_1\, q(t) - \mu_2\, w_N(t)), \qquad \mu_1, \mu_2 > 0,$$

where $q(t) = \Phi(t;\kappa)/(\rho\, U(t))$ is the inflow-to-throughput ratio. High κ raises q, which raises μ through $\mu_1 > 0$; the resulting wage depression further raises μ through $\mu_2 > 0$.

The minimal model abstracts from productivity gains associated with automation, focusing on a transition phase in which substitution effects dominate and labor reallocation drives wage dynamics.

Welfare

The planner values output gains from automation through a bounded increasing function $V_Y(\kappa)$ and bears social costs $\chi_U, \chi_D > 0$ with $\chi_D > \chi_U$ for unemployment and permanent exit:

$$W(\kappa) = V_Y(\kappa) - \chi_U \int_0^\infty e^{-\beta t} U(t;\kappa)\, dt - \chi_D \int_0^\infty e^{-\beta t} D(t;\kappa)\, dt, \quad \beta > 0.$$

Four results

M1 (Unemployment hump). $U(t;\kappa)$ is hump-shaped. Peak unemployment $U_{\max}$ is strictly increasing in κ and $\bar{\imath}$, and strictly decreasing in ρ. Higher

adoption speed raises the inflow $\Phi_{\max} = \bar{\iota}\kappa/4$ without raising the retraining throughput ρ, so the queue grows faster.

M2 (Permanent exit). $D(\infty;\kappa)$ is strictly increasing in κ (for $\mu_1 > 0$ or $\mu_2 > 0$) and strictly decreasing in ρ. Faster adoption permanently reduces labor force participation.

M3 (Wage and labor-intensity crossing). Under fast adoption, non-routine labor intensity $L_N(t)$ is higher early (retrained workers arrive faster) but lower in the long run (permanent exit depletes the retrained stock). Wages cross in the mirror-image direction: lower during the transition, higher in the long run. Fast adoption can therefore raise long-run wages for workers who remain in the non-routine sector while permanently reducing participation.

The model implies three distinct distributional margins. First, faster adoption increases inequality between employed and inactive workers by raising long-run wages for a smaller active workforce while expanding permanent exit. Second, it generates implicit cohort inequality within non-routine workers, as early entrants face wage compression while later survivors benefit from scarcity. Third, it can reduce aggregate labor income—and thus the labor share—even as wages rise, due to declining participation.

M4 (Interior optimal speed). $W(\kappa)$ is maximized at an interior $\kappa^* \in (0, \infty)$. Because the automation frontier is fixed, faster adoption re-times but does not expand total productivity gains, implying diminishing marginal returns to speed as diffusion approaches completion. Output gains $V_Y(\kappa)$ saturate as the logistic completes; transition costs rise with peak displacement and are amplified by the irreversibility premium $\chi_D > \chi_U$. The market speed $\kappa^{\text{mkt}} > \kappa^*$ because firms do not internalize congestion externalities or permanent exit.

What the full model adds

The minimal model above is sufficient to establish the qualitative mechanism. The full model in Section 4 adds three things that matter for quantification and credibility: *(i)* it derives the wage equation from a CES task-based production structure, which pins σ to task substitution elasticities measured in the micro literature and makes the wage response a theorem rather than an assumption; *(ii)* it derives $V_Y(\kappa)$ from production primitives—the logistic diffusion of a_t interacting with the CES aggregator—rather than asserting boundedness; and *(iii)* it

generates calibration targets in Appendix D from observable PIAAC routine task shares and OECD ALMP spending, providing a basis for quantitative evaluation of the welfare wedge $\kappa^{\text{mkt}} - \kappa^*$. None of those additions alter the mechanism: for a fixed long-run automation envelope, adoption speed interacts with finite retraining capacity to shape congestion, discouragement, and long-run participation.

4 The Model

The economy transitions toward full automation of its routine task sector. The long-run automation threshold $\bar{\imath} \in (0, 1)$ is taken as given—the share of tasks eventually automated, determined by technology and task structure (the Acemoglu–Restrepo envelope). We model the speed κ at which this transition occurs and its labor market consequences. The working-age population is normalized to unity.

4.1 Technology and Displacement Flow

AI adoption cost falls over time as firms learn, regulation adapts, and technology diffuses. The cumulative share of routine tasks automated is:

$$A(t) = \bar{\imath} \cdot a_t \tag{1}$$

where $a_t \in [0, 1]$ follows logistic diffusion:

$$\frac{da_t}{dt} = \kappa \cdot a_t(1 - a_t), \qquad a_0 \in (0, 1) \tag{2}$$

with solution $a_t = 1/\big(1 + [(1 - a_0)/a_0]\, e^{-\kappa t}\big)$. Higher κ compresses the timeline without changing $A(\infty) = \bar{\imath}$: same destination, different transition speed.

The flow of newly displaced routine workers is:

$$\Phi(t; \kappa) = \bar{\imath} \cdot \kappa \cdot a_t(1 - a_t) \tag{3}$$

$\Phi(t; \kappa)$ peaks at the logistic inflection $a_t = 1/2$, where $\Phi_{\max} = \bar{\imath}\kappa/4$. Peak displacement is linear in κ: doubling adoption speed doubles the worst-case displacement flow. Total displacement satisfies $\int_0^\infty \Phi(t; \kappa)\, dt = \bar{\imath}$ for all κ—same total displaced workers regardless of speed.

Notation. The mapping from task measure to worker headcount is one-to-one: each routine task employs exactly one worker, so $\bar{\imath}$ represents both the automat-

able task share and the corresponding worker share. Tasks $i > \bar{\imath}$ require human judgment not substitutable by AI on the horizon studied.

4.2 Labor Market: Four States

Workers occupy one of four mutually exclusive states at each instant:

State	Description
$L^R(t)$	Employed in remaining (not yet automated) routine tasks
$L^N(t)$	Employed in non-routine tasks — grows as retrained workers arrive
$U(t)$	Unemployed: displaced, searching / retraining
$D(t)$	Permanently discouraged: out of labor force (absorbing state)

Accounting identities. Let $R(t)$ denote the cumulative stock of workers who have retrained and moved from the routine base into L^N. Since all discouragement originates from displaced routine workers (Assumption 1), the routine base satisfies:

$$L^R(t) + U(t) + D(t) + R(t) = \bar{\imath} \tag{4}$$

Non-routine employment:

$$L^N(t) = (1 - \bar{\imath}) + R(t) \tag{5}$$

Adding (4) and (5): $L^R + L^N + U + D = \bar{\imath} + (1 - \bar{\imath}) = 1$. Population identity holds exactly.

Assumption 1 (Source of discouragement)**.** *$D(t)$ consists entirely of workers displaced from the routine sector. Non-routine incumbent workers do not exit the labor force on the horizon studied. This is maintained throughout.*

4.3 Flow Equations

Three ODEs govern the transition dynamics:

$$\frac{dU}{dt} = \Phi(t;\kappa) - \left[\rho(Z) + \mu(q(t), w^N(t))\right] \cdot U(t) \tag{6}$$

$$\frac{dL^N}{dt} = \frac{dR}{dt} = \rho(Z) \cdot U(t) \tag{7}$$

$$\frac{dD}{dt} = \mu(q(t), w^N(t)) \cdot U(t) \tag{8}$$

where $\rho(Z) > 0$ is the reallocation hazard rate—the probability per unit time that a worker in U successfully retrains and enters L^N—depending on institutional factors Z (ALMP spending, employment protection, education, mobility); and $\mu(\cdot) > 0$ is the discouragement hazard defined in Section 2.4.

Competing drains on U. The two hazards $\rho(Z)$ and $\mu(q, w^N)$ are competing outflows from the same stock $U(t)$. This competition is the central mechanism of the model. Even with $\rho(Z)$ fixed institutionally, a rise in μ depletes U faster, leaving fewer workers available to retrain. Formally:

$$\frac{dR}{dt} = \rho(Z) \cdot U(t) \quad \Rightarrow \quad R(t) = \rho(Z) \int_0^t U(s)\, ds \tag{9}$$

So anything that reduces $U(s)$—including accelerated exit via higher μ—reduces cumulative retraining $R(t)$ even though $\rho(Z)$ is unchanged. Wage depression therefore slows retraining accumulation *through the exit margin*, without requiring ρ to depend on w^N. The feedback loop is:

$$w^N \downarrow \Rightarrow\ \mu \uparrow \Rightarrow\ U \downarrow \Rightarrow\ \rho(Z) \cdot U \downarrow \Rightarrow\ R(t) \downarrow \Rightarrow\ L^N \text{ grows slower}$$

Flow consistency. Total outflow from U is $[\rho+\mu] \cdot U = dL^N/dt + dD/dt$. Every displaced worker either retrains ($\Rightarrow L^N$) or exits permanently ($\Rightarrow D$). No worker is created or destroyed.

4.4 Non-Routine Labor Market: Elastic Demand, Inelastic Supply, Demand-Determined Wage

Firm-side demand. Output is produced via a CES aggregator over tasks $i \in [0, 1]$ (full expression in Appendix A). Within the non-routine task region $i \in (\bar{\imath}, 1]$, firms choose labor intensity $N(i, t)$ to maximize profit at the going wage $w^N(t)$. Under symmetric task assignment $N(i, t) = L^N(t)/(1 - \bar{\imath})$ for all i, the firm's first-order condition equates the non-routine wage to the marginal product of labor:

$$w^N(t) = A_N^{1-1/\sigma} \cdot \left(\frac{Y_t}{L^N(t)} \right)^{1/\sigma} \cdot (1 - \bar{\imath})^{1/\sigma} \tag{10}$$

Inverting (10) gives the labor demand curve:

$$L^{N,d}(w^N(t)) = A_N^{\sigma-1} \cdot (1 - \bar{\imath}) \cdot \left(\frac{Y_t}{w^N(t)} \right)^{\sigma} \tag{11}$$

Demand is *elastic*: the elasticity of labor demand with respect to w^N is $-\sigma < -1$ (since $\sigma > 1$). Firms willingly absorb additional non-routine workers at lower wages by increasing labor intensity within each fixed task. Lower wages make it profitable to staff each task more intensively. This is the mechanism by which displaced routine workers find room in the non-routine sector—not through the creation of new tasks, but through factor price adjustment within a fixed task measure $(1 - \bar{\imath})$.

Remark 1 (Comparison with Acemoglu–Restrepo)**.** *The present model abstracts from endogenous task creation, treating the non-routine task measure* $(1 - \bar{\imath})$ *as fixed on the horizon studied. This is consistent with the empirical observation that reinstatement of new task categories operates on longer horizons than the displacement effects of automation waves [Acemoglu and Restrepo, 2022]. Reallocation therefore occurs entirely through factor price adjustment within the fixed task measure, not through task expansion. Section 6 introduces a second diffusion process for task creation at speed* κ_T *and shows that the AR long-run benchmark is recovered as* $\kappa_T \to \infty$*.*

Assumption 2. *(Inelastic retraining success conditional on activity). Displaced routine workers do not search for routine jobs, since automated routine tasks do not reappear on the horizon studied. Conditional on remaining active in retraining/search* U*, successful transitions into the non-routine sector occur at rate* $\rho(Z)\,U(t)$*, where* $\rho(Z)$ *is determined by institutional capacity* Z*—including training infrastructure, ALMP programs, and educational quality—and is independent of the current non-routine wage* $w^N(t)$*. Wages affect reallocation only through the exit margin* $\mu(q(t), w^N(t))$*, by changing the attractiveness of remaining active rather than the conditional success rate of those who stay.*

Behavioral foundation: stay-or-go reassessment. Workers in U are *active* job-seekers: remaining in U means continuing to search and retrain. There is no passive waiting — a worker either stays active and retrains, or exits permanently to D. The two hazards govern *distinct margins*:

$\mu(q, w^N)$ governs the *stay-or-go* decision. At each instant, the worker reassesses whether continued activity is worthwhile given $V_U(t)$ relative to V_D, as formalized in (19)–(20). Congestion raises the flow cost $-c_1 q$ of staying active; lower wages reduce the payoff $V_N(w^N)$ to eventual success. Both forces lower V_U relative to V_D, raising the exit hazard.

$\rho(Z)$ governs *retraining success conditional on staying active*. A worker who remains in U succeeds at the institutionally-determined rate $\rho(Z)$, reflecting training infrastructure, ALMP capacity, and slot availability — not individual incentives or wages. Just as a doctoral student's departmental pass rate is set by the institution regardless of how attractive the job market looks, $\rho(Z)$ is outside any individual worker's control.

The two parameters therefore operate on distinct objects with no behavioral inconsistency: μ determines who leaves U via the value comparison (20); ρ determines who succeeds among those who stay. Assumption 2 is the direct consequence: the flow into L^N is $\rho(Z) \cdot U(t)$, independent of w^N, because wages govern the stay-or-go decision through V_U but not the success rate of those who remain active. Thus, at each instant, the stock of non-routine labor is predetermined by past retraining flows. Adjustment to wage changes operates through participation and retraining dynamics rather than contemporaneous labor supply responses. [2]

The supply of non-routine labor at each instant is:

$$L^N(t) = (1 - \bar{\imath}) + R(t) \tag{12}$$

This supply is predetermined at each t: it is the outcome of all prior retraining flows, governed by the ODE system (6)–(8).

Remark. Firm-level retraining or internal redeployment would reduce net inflows into U or increase effective capacity Z. We abstract from this margin to isolate system-wide congestion in retraining.

Market clearing. At each instant, inelastic supply $L^N(t)$ meets elastic demand $L^{N,d}(w^N(t))$. Since demand is continuous and strictly decreasing in w^N, and $L^{N,d}(0) = \infty$, $L^{N,d}(\infty) = 0$, there exists a unique market-clearing wage $w^N(t)$ satisfying:

$$L^{N,d}(w^N(t)) = L^N(t) \tag{13}$$

at every t. Wages adjust instantaneously. Evaluating (10) at the market-clearing quantity and using $w^N(0) = w_0^N$ at $L^N(0) = 1 - \bar{\imath}$ yields the exact equilibrium wage path:

$$w^N(t) = w_0^N \cdot \left[\frac{1 - \bar{\imath}}{L^N(t)}\right]^{1/\sigma} \cdot \left[\frac{Y_t}{Y_0}\right]^{1/\sigma} \tag{14}$$

[2] We abstract from wage-induced exit among incumbent non-routine workers. In principle, sufficiently low w^N could also reduce participation on the non-routine side. Incorporating this margin would add an additional channel of labor supply adjustment but would likely reinforce, rather than overturn, the welfare costs of rapid adoption emphasized here.

Equation (14) makes clear that faster adoption has two opposing effects on non-routine wages. Differentiating logarithmically with respect to adoption speed κ,

$$\frac{\partial \ln w^N(t)}{\partial \kappa} = \frac{1}{\sigma}\left[\frac{\partial \ln Y_t}{\partial \kappa} - \frac{\partial \ln L^N(t)}{\partial \kappa}\right].$$

The first term is a demand effect: faster automation raises output and shifts demand for non-routine labor outward. The second is a supply effect: faster displacement and retraining increase non-routine labor supply and depress wages. The sign of $\partial w^N(t)/\partial \kappa$ is therefore ambiguous in general. This decomposition highlights that non-routine workers benefit from productivity gains associated with automation, mitigating the compression of non-routine wages due to the supply of retrained workers.

For the baseline analysis, we adopt the reduced-form approximation in which the labor-supply effect dominates during the transition window, abstracting from the output-growth term to isolate the wage effects of labor-supply expansion through retraining:

$$\boxed{w^N(t) \approx w_0^N \cdot \left[\frac{1-\bar{\imath}}{L^N(t)}\right]^{1/\sigma}} \tag{15}$$

Thus, the baseline abstracts from technology-induced wage gains for non-routine workers to isolate the transitional wage compression generated by congestion in labor reallocation.

Assumption 3. *(Labor-supply effect dominates during transition). The AI capital share of output $s_K(t)$ is sufficiently small during the transition that the output-growth offset in (14) is negligible relative to the labor-supply depression term. This corresponds to an early phase of adoption in which substitution effects dominate over productivity-enhancing complementarities.*

Derivation of the approximation condition. From the CES system (Appendix A), the elasticity of Y_t with respect to a_t is:

$$\frac{\partial \ln Y_t}{\partial \ln a_t} = s_K(t) \cdot \frac{\sigma - 1}{\sigma}$$

where $s_K(t)$ is the AI capital income share. The labor-supply depression term has elasticity $1/\sigma$ with respect to L^N. The approximation $Y_t/Y_0 \approx 1$ is valid when the output-growth contribution to w^N is small relative to the labor-supply depression

contribution, i.e.:

$$s_K(t) \cdot (\sigma - 1) \ll 1 \tag{16}$$

During early transition $s_K(t) \approx s_{K,0}$, so the condition reduces to $s_{K,0}(\sigma - 1) < 1$, giving an explicit upper bound:

$$\sigma < \bar{\sigma} \equiv 1 + \frac{1}{s_{K,0}} \tag{17}$$

With $s_{K,0} \approx 0.05$ (AI capital initially a minor share of output, consistent with Appendix D calibration targets), this yields $\bar{\sigma} \approx 21$, well above the calibrated range $\sigma \in [1.5, 4]$. Assumption 3 is therefore satisfied throughout the relevant parameter space.

Remark 2 (Conservative direction)**.** *The output-growth term in* (14) *attenuates wage depression. Dropping it overstates* w^N *depression, overstates* $D(\infty)$*, and therefore overstates the optimal policy wedge* $\kappa^{mkt} - \kappa^*$*. The approximate model yields an* upper bound *on the case for intervention: the exact model* (14) *would support a smaller but still positive wedge. Policy recommendations based on* (15) *are therefore conservative in the direction of recommending intervention.*

Equation (15) is a *result*—the market-clearing wage—not an assumption. Since $L^N(t) > 1 - \bar{\imath}$ for all $t > 0$ (retraining adds workers above the initial base), $w^N(t) < w_0^N$: the non-routine wage falls monotonically below its initial level as retrained workers accumulate.

Economic interpretation. The non-routine labor market is too efficient in a dynamic sense: it clears at every instant by depressing wages, which is privately optimal for firms (they absorb cheap labor) but socially costly because the same price signal that equilibrates today's market discourages tomorrow's retraining. Each successfully retrained worker imposes a negative wage externality on all workers still in U, raising their discouragement hazard μ via $\mu_2 < 0$. No individual worker internalizes this externality.

4.5 Congestion Index and Discouragement Hazard

Congestion index. $\rho(Z)$ is best understood as a *system-level retraining capacity*: the total throughput of the economy's retraining infrastructure (training centers, ALMP slots, employer absorption capacity) per unit of the existing queue,

as determined by institutional factors Z. Firm-sponsored retraining can also mitigate these frictions; we return to this margin in the conclusion.

The total retraining slots available per unit time is $\rho(Z) \cdot U(t)$. When new displacement $\Phi(t;\kappa)$ arrives faster than the system can process it, slots are rationed and individual success rates fall. The congestion index captures this utilization rate:

$$q(t) = \frac{\Phi(t;\kappa)}{\rho(Z) \cdot \max\{U(t), \varepsilon\}} \tag{18}$$

where $\varepsilon > 0$ prevents mechanical singularity when $U \to 0$ at the end of the transition. $q(t) > 1$ means new displacement inflow exceeds retraining throughput: the system is congested, slots are over-subscribed, and individual expected waiting times rise. Workers in U observe $q(t)$ as a signal of their individual retraining prospects and update their exit decision accordingly via $\mu_1 > 0$.

Robustness to ε. As $U(t) \to \varepsilon$ at the end of the transition, $\Phi(t;\kappa) \to 0$ (displacement is complete), so $q(t) \to \Phi/(\rho\varepsilon) \to 0$. Under the linear specification (22), $\mu \to \mu_0 - \alpha_2 w_\infty^N$, and the discouragement flow $\mu \cdot U \to (\mu_0 - \alpha_2 w_\infty^N)\varepsilon \to 0$ as $\varepsilon \to 0$. The contribution of the drainage phase to $D(\infty)$ therefore vanishes with ε, and all results are robust to $\varepsilon \to 0$.

Dimensional note. Φ is a flow (mass/time) and ρ is a hazard (1/time), so Φ/ρ has units of mass. Dividing by U gives the dimensionless $q = \Phi/(\rho \cdot U)$, the correct capacity utilization measure.

Note: $q > 1$ does not imply $dU/dt > 0$, because the discouragement outflow μU also drains U—the queue can shrink even under pipeline congestion.

Discouragement hazard.

Optimal stopping foundation. We microfound μ as the endogenous exit rate from an optimal stopping problem. Let $V_U(t)$ be the value of being active in U at time t, $V_D = b/\beta$ the value of permanent exit (constant outside option b, e.g. home production or informal work), and $V_N(w^N(t)) = u_N(w^N(t))/\beta$ the value of entering non-routine employment at the market-clearing wage, with $u_N' > 0$.

The Hamilton–Jacobi–Bellman equation for an active worker is:

$$\beta V_U(t) = u_U(q(t)) + \rho(Z)\big[V_N(w^N(t)) - V_U(t)\big] + \dot{V}_U(t) \tag{19}$$

where $u_U(q(t)) = -c_0 - c_1 q(t)$ is the flow utility of staying active, with $c_0, c_1 >$

0. The term $-c_1 q(t)$ captures the *flow cost of congestion*: when the retraining system is overloaded, workers face longer queues, greater search effort, and higher psychic costs of continued non-employment. Congestion therefore enters through the cost of staying active, not through the success technology $\rho(Z)$. This is the key separation: $\rho(Z)$ is the institutional throughput rate conditional on active search; $q(t)$ governs how costly active search is.

The worker exits to D when $V_D \geq V_U(t)$. In the pure threshold model this generates bang-bang exit, producing discontinuities in aggregate dynamics. To obtain a smooth aggregate hazard we follow the standard approach and introduce idiosyncratic taste shocks: at each instant, worker i draws a preference shock ε_i for continued activity (i.i.d. extreme value), smoothing the exit decision into a logistic probability. The resulting aggregate exit hazard is:

$$\mu(t) = \bar{\mu} \cdot \Lambda\big(V_D - V_U(t)\big) \tag{20}$$

where $\Lambda(x) = 1/(1 + e^{-x})$ is the logistic function and $\bar{\mu} > 0$ is the maximum exit intensity. Since Λ is smooth and bounded, $\mu(t) \in (0, \bar{\mu})$ at all times — nonnegativity and boundedness hold automatically without auxiliary restrictions.

Signed partial derivatives. From (19) and (20):

- $\partial\mu/\partial w^N < 0$: higher w^N raises V_N and therefore V_U via the success payoff $\rho(Z)[V_N - V_U]$; this narrows $V_D - V_U$ and lowers μ.
- $\partial\mu/\partial q > 0$: higher q lowers $u_U(q)$ and therefore V_U; this widens $V_D - V_U$ and raises μ.

Collecting these into reduced form:

$$\mu = \mu(q(t),\, w^N(t)), \qquad \mu_1 \equiv \frac{\partial \mu}{\partial q} > 0, \quad \mu_2 \equiv \frac{\partial \mu}{\partial w^N} < 0 \tag{21}$$

Baseline parametric specification. For the baseline analysis and calibration we adopt the linear specification:

$$\mu(q(t), w^N(t)) = \mu_0 + \alpha_1\, q(t) - \alpha_2\, w^N(t) \tag{22}$$

with $\mu_0, \alpha_1, \alpha_2 > 0$. The hazard is truncated below at zero, $\mu = \max(0,\, \mu_0 + \alpha_1 q - \alpha_2 w^N)$, ensuring non-negativity for all t. When w^N is near its initial value $w_0^N = 1$ and q is small, $\mu \approx \mu_0 > 0$; as w^N falls and q rises during the transition,

μ increases above the baseline. The first-order Taylor expansion of (20) around the baseline value $\bar{V} \equiv V_D - V_U(0)$:

$$\mu \approx \bar{\mu} \cdot \Lambda(\bar{V}) + \bar{\mu} \cdot \Lambda'(\bar{V}) \cdot \frac{\partial(V_D - V_U)}{\partial q} \Delta q + \bar{\mu} \cdot \Lambda'(\bar{V}) \cdot \frac{\partial(V_D - V_U)}{\partial w^N} \Delta w^N$$

which maps to (22) with $\mu_0 = \bar{\mu}\Lambda(\bar{V})$, $\alpha_1 = \bar{\mu}\Lambda'(\bar{V}) \cdot c_1/(\beta + \rho)$, and $\alpha_2 = \bar{\mu}\Lambda'(\bar{V}) \cdot u_N'(w_0^N)/(\beta + \rho)$. Linearity is thus a local approximation to the logit hazard, valid when Δq and Δw^N are not too large relative to the baseline — a condition satisfied in the calibrated parameter range.

The linear specification is Lipschitz with constant $L = \alpha_1 \|\partial q/\partial U\|_\infty + \alpha_2 \|\partial w^N/\partial L^N\|_\infty$, both finite for $U \geq \varepsilon > 0$ and $L^N \geq 1 - \bar{\imath} > 0$. More fundamentally, the logit hazard (20) is globally Lipschitz since Λ' is bounded: $|\Lambda'(x)| \leq 1/4$ for all x. Assumption 4 is therefore verified directly from the smoothness of the logistic function, without relying on the linear approximation.

Two forces drive μ upward under fast κ:

- **Queue pressure** ($\mu_1 > 0$). High $q(t)$ signals long expected retraining queues. Workers rationally update downward on re-employment prospects.
- **Wage depression** ($\mu_2 < 0$). As L^N grows, w^N falls via (15). Lower available wages make retraining less attractive relative to permanent exit.

κ affects μ only through $q(t)$ and $w^N(t)$—not directly. The two channels reinforce each other: high κ raises Φ, increasing both q and the rate of L^N accumulation, which depresses w^N. Discouragement is therefore superlinear in κ, driving the interior optimum.

Full feedback loop. Combining the labor market clearing condition (15) with the competing hazards structure (6)–(8) and the discouragement hazard (21), the complete feedback mechanism is:

$$\kappa \uparrow \Rightarrow \Phi \uparrow \Rightarrow \begin{cases} q \uparrow & \Rightarrow \mu \uparrow \text{ (queue channel)} \\ R(t) \uparrow \Rightarrow L^N \uparrow \Rightarrow w^N \downarrow & \Rightarrow \mu \uparrow \text{ (wage channel)} \end{cases} \Rightarrow U \downarrow \Rightarrow \begin{cases} \rho(Z) \cdot U \downarrow & \text{(re} \\ D(\infty) \uparrow & \text{(pe} \end{cases}$$

The stabilizing force operates entirely through the price mechanism. Higher κ clears the non-routine labor market faster and at lower wages, which is privately optimal for firms but generates excessive permanent exit in aggregate.

5 Welfare

Social welfare trades off output gains from faster automation against two transition costs:

$$W(\kappa) = \int_0^\infty e^{-\beta t}\big[Y(t;\kappa) - \chi^U \cdot U(t;\kappa) - \chi^D \cdot D(t;\kappa)\big]\, dt \tag{23}$$

where $Y(t;\kappa)$ is aggregate output from the CES system (Appendix A); $\chi^U > 0$ is the flow disutility of unemployment (foregone wages, search costs, psychic costs); and $\chi^D > \chi^U$ is the permanent disutility of labor force exit (irreversibility premium: lost human capital, health costs of non-participation, loss of social insurance).[3]

The analysis requires only that faster adoption yields bounded and eventually diminishing private surplus gains — specifically that $V_Y(\kappa)$ is concave and bounded above by $Y(\bar{\imath})/\beta$, as established in Appendix B. No assumption of sustained productivity growth or output explosion is needed: the marginal benefit of speed is positive but diminishing, and it is this saturation that generates the interior optimum κ^*.

Wage depression enters through $Y(t)$: as w^N falls, income shifts from high-MPC routine/retrained workers toward capital owners, reducing aggregate consumption relative to output.

The first-order condition for κ^*:

$$\int e^{-\beta t}\frac{\partial Y}{\partial \kappa}\, dt = \chi^U \int e^{-\beta t}\frac{\partial U}{\partial \kappa}\, dt + \chi^D \int e^{-\beta t}\frac{\partial D}{\partial \kappa}\, dt \tag{24}$$

Left side: marginal social benefit of faster adoption (earlier productivity gains). Right side: marginal social cost of increased unemployment and permanent discouragement.

Market equilibrium anchor. The market speed κ^{mkt} is determined by firm-level adoption decisions: each firm adopts when the private NPV of automation (productivity gain minus adoption cost) is positive, internalizing neither the congestion externality on other workers' job-finding rates nor the irreversibility of discouraged exit. The firm's private FOC omits the χ^U and χ^D terms on the

[3]There is no double-counting between the $Y(t)$ and $\chi^D D(t)$ terms. $Y(t)$ captures aggregate output, which falls as workers exit to D via reduced labor input — this is the *pecuniary* cost of exit already embedded in the output measure. χ^D captures exclusively the *non-pecuniary* irreversibility premium: the psychic cost of permanent non-participation, health deterioration, loss of social insurance eligibility, and human capital depreciation that are not reflected in market output. The two terms therefore measure distinct welfare components with no overlap.

right-hand side of (24). Since both terms are positive and $\partial U/\partial\kappa > 0$, $\partial D/\partial\kappa > 0$ during the transition, the private marginal benefit exceeds the social marginal benefit at any κ, yielding $\kappa^{\text{mkt}} > \kappa^*$ by single-crossing.

What firms are choosing. κ is not a primitive of nature but the aggregate outcome of decentralized firm decisions. Each firm chooses the intensity of deployment effort: investment in AI integration, workflow reorganization, compliance infrastructure, and staff training for the technology itself. These choices determine how fast a_t evolves at the firm level; aggregated across firms they generate κ^{mkt}. Firms optimize private NPV, ignoring two externalities. First, each firm's displacement of routine workers congests the shared retraining pipeline, reducing other displaced workers' retraining success rates. Second, the irreversibility of discouraged exit imposes costs that are not priced into the adoption decision. Wage depression amplifies the discouragement channel but is partially internalized through the reduction in cost savings from automating additional tasks as w^N falls. Together these mechanisms ensure $\kappa^{\text{mkt}} > \kappa^*$.

Which externality is structurally necessary. The interior optimum requires that permanent exit be socially costlier than temporary unemployment ($\chi^D > \chi^U$, Assumption 2); without this irreversibility premium the cost side of the planner's FOC is weakened and the interior solution may not exist. Of the two remaining channels, the *wage depression* amplifies the discouragement hazard ($\mu_2 < 0$) rather than acting as an independent externality: it raises μ conditional on congestion, strengthening but not substituting for the irreversibility mechanism. The congestion externality alone ($\mu_1 > 0$, $\chi^D = \chi^U$) generates excess unemployment but not permanent exit, which weakens the cost side of (24) and may not be sufficient to overturn the private adoption incentive. The irreversibility premium makes permanent exit disproportionately costly; the wage channel ensures that fast adoption generates permanent exit through the discouragement mechanism. Both are needed.

Partial internalization. One natural objection is that firms partially internalize the wage externality: falling non-routine wages reduce the cost advantage of automation relative to labor, dampening private adoption incentives. This is correct in principle — lower w^N reduces the firm's cost savings from automating additional tasks, creating a second-order brake on κ^{mkt}. However, the dominant driver of private adoption is the productivity gain from automating routine tasks,

not the cost differential with non-routine labor. For $AK/r > AR/w^R$ (the adoption condition in Appendix A), the private NPV of automation is positive and large, so the wage feedback is a second-order correction that does not eliminate the wedge $\kappa^{\text{mkt}} > \kappa^*$. The single-crossing argument holds for any $\chi^D > \bar{\chi}^D$ (Assumption 3), regardless of the partial internalization effect.

Knife-edge cases and robustness. Four parametric configurations collapse the interior optimum to a corner or make κ^* degenerate. (i) $\chi^D = \chi^U$: permanent exit carries the same flow cost as unemployment; the irreversibility premium vanishes and κ^* may equal κ^{mkt} or be indeterminate. (ii) $\mu_2 = 0$: wages do not affect discouragement; the wage depression channel is severed and $D(\infty)$ becomes less sensitive to κ, weakening the cost curvature needed for an interior solution. (iii) $\kappa_T \to \infty$ (Section 7): instantaneous task creation absorbs displaced workers before discouragement accumulates; the Acemoglu–Restrepo benchmark is recovered and $D(\infty) \to 0$. (iv) $\sigma \to \infty$: the wage depression vanishes (perfectly elastic non-routine labor demand); discouragement operates only through congestion, reducing but not eliminating the wedge. Assumption 3 rules out cases (i) and (ii) by requiring $\chi^D > \bar{\chi}^D$.

6 Main Results

Proposition 1 (Unemployment Hump)**.** *Under displacement flow* (3)*, $U(t)$ follows a hump-shaped path. Peak unemployment $U_{\max}$ is:* (i) *strictly increasing in κ;* (ii) *strictly increasing in $\bar{\iota}$;* (iii) *strictly decreasing in $\rho(Z)$.*

Proof sketch. From (6), $dU/dt > 0$ iff $\Phi(t;\kappa) > [\rho + \mu] \cdot U$. Since $\Phi(t;\kappa)$ is hump-shaped (logistic derivative) and $[\rho+\mu]U$ starts near zero, $dU/dt > 0$ initially. At peak U, $dU/dt = 0$, giving $U_{\max} = \Phi(t_{\text{peak}};\kappa)/[\rho + \mu(q, w^N)]$. Since $\Phi_{\max} = \bar{\iota}\kappa/4$ is linear in κ and $\bar{\iota}$, and $[\rho+\mu] > 0$ is bounded away from zero, $U_{\max}$ is increasing in κ and $\bar{\iota}$. Decreasing in $\rho(Z)$ follows from the denominator. □

Proposition 2 (Permanent Labor Force Exit)**.** *The long-run discouraged stock $D(\infty) \in (0, \bar{\iota})$. Under Assumption 1 and $\mu_1 > 0$, $\mu_2 < 0$:*

(i) *$D(\infty)$ is strictly increasing in κ*

(ii) *$D(\infty)$ is strictly decreasing in $\rho(Z)$*

(iii) *$D(\infty)$ is strictly increasing as σ decreases (stronger wage feedback)*

Remark 3. *No closed-form expression for $D(\infty)$ is asserted in general, since μ is endogenous and time-varying. A constant-hazard approximation appears in Appendix C.*

Proof sketch. $D(\infty) = \int_0^\infty \mu(q(s), w^N(s)) \cdot U(s)\, ds$.

Part (i): Higher κ raises $\Phi(t;\kappa)$ for $t \leq t^*$ (the acceleration phase), increasing $U(s)$ and $q(s) = \Phi(s)/(\rho \cdot \max\{U(s), \varepsilon\})$, raising μ via $\mu_1 > 0$. Simultaneously, faster L^N growth depresses w^N via (15), raising μ via $\mu_2 < 0$. Crucially, higher μ depletes U faster, reducing $\rho(Z) \cdot U$ and therefore slowing $R(t)$ accumulation, which further depresses w^N—a self-reinforcing loop. Both direct effects increase the integrand on a positive-measure set of s, so $D(\infty)$ rises.

Parts (ii)–(iii): Analogously. Higher ρ reduces q and slows L^N accumulation; lower σ steepens the w^N response to L^N growth. □

Proposition 3 (Non-Routine Wage Depression)**.** *Let T be the end of the active displacement phase (first time $\Phi(T;\kappa) \leq \varepsilon$ for small $\varepsilon > 0$). Then:*

(i) *For all $t \in (0, T]$: $\partial w^N(t)/\partial\kappa < 0$. The non-routine wage is strictly lower under faster adoption during the transition window.*

(ii) *For $t > T$: $\Phi(t;\kappa) \approx 0$ and $U(t)$ drains toward zero. From the accounting identity (4) with all terms non-negative, $R(\infty) = \bar{\imath} - D(\infty) - L^R(\infty) \leq \bar{\imath}$, so $R(\infty)$ is finite and well-defined. $L^N(t)$ therefore stabilizes at $L^N_\infty = (1 - \bar{\imath}) + R(\infty) \leq 1$, and $w^N(t)$ converges monotonically to:*

$$w^N_\infty = w^N_0 \cdot \left[\frac{1-\bar{\imath}}{L^N_\infty}\right]^{1/\sigma} < w^N_0$$

Since $R(\infty) = \bar{\imath} - D(\infty)$ (with $L^R(\infty) = 0$ at full automation) and $D(\infty)$ is increasing in κ (Proposition 2), $R(\infty)$ is decreasing *in κ: faster adoption sends more workers permanently to D and fewer to L^N. Therefore L^N_∞ is smaller and w^N_∞ is* higher *under faster adoption in the long run. The long-run wage level is thus non-monotone in κ relative to the transition path: wages fall more steeply* during *the transition under high κ (part i), but the long-run level w^N_∞ is partially restored because fewer workers complete retraining. This does not overturn the welfare result: the discounted transition-window depression dominates by part (iii).*

(iii) *The welfare-relevant cumulative discounted wage depression is unambiguously increasing in* κ*:*

$$\Delta W^N(\kappa) \equiv \int_0^\infty e^{-\beta t}\big[w_0^N - w^N(t;\kappa)\big]\,dt \quad \textit{is strictly increasing in } \kappa$$

Proof of (i) and (iii). From (15): $\partial w^N/\partial\kappa = -(1/\sigma)(w^N/L^N)\cdot\partial L^N/\partial\kappa$. It suffices to show $\partial L^N(t)/\partial\kappa > 0$ for $t \in (0,T]$. From (5) and (7): $L^N(t) = (1-\bar{\imath}) + \rho(Z)\int_0^t U(s)\,ds$, so $\partial L^N/\partial\kappa = \rho(Z)\int_0^t \partial U(s)/\partial\kappa\,ds$. Using the Duhamel representation with approximately constant hazard rates (conservative; $\mu_1 > 0$ strengthens the result), this integral is dominated by the contribution from $s \le t^*$ where $\partial\Phi/\partial\kappa > 0$, giving a positive result for all $t \le T$. Hence $\partial L^N/\partial\kappa > 0$ and $\partial w^N/\partial\kappa < 0$ on $(0,T]$.

Part (iii): Since $\partial w^N/\partial\kappa < 0$ for $t \in (0,T]$ and $e^{-\beta t}$ is highest during the transition window, $\partial\Delta W^N/\partial\kappa > 0$. The post-$T$ recovery is discounted more heavily and does not overturn the net effect for any $\beta > 0$. □

Remark 4 (The crossing result)**.** *Faster adoption increases non-routine labor intensity and depresses wages during the transition, but raises long-run wages by reducing the stock of successfully retrained workers through higher permanent exit. Formally:* $L^N(t;\kappa_H) > L^N(t;\kappa_L)$ *for* $t \in (0,t_c)$ *and* $L^N(t;\kappa_H) < L^N(t;\kappa_L)$ *for* $t > t_c$*, where* t_c *is the crossing time. The implied wage crossing is* $w^N(t;\kappa_H) < w^N(t;\kappa_L)$ *for* $t < t_c$ *and* $w^N(t;\kappa_H) > w^N(t;\kappa_L)$ *for* $t > t_c$*. In the calibration of Figure 1 (*$\alpha_2 = 0.15$*,* $\sigma = 1.5$*),* $t_c \approx 33$ *years, the long-run* L_N *gap is 8 percentage points (*0.835 *vs.* 0.915*), and wages end 5 pp higher under fast adoption (*0.846 *vs.* 0.796*). High-speed transitions can therefore end with higher wages for surviving non-routine workers alongside permanently lower labor force participation — a distributional outcome that wage-level statistics alone would mischaracterize as beneficial.*

Proposition 4 (Interior Optimal Adoption Speed)**.** *Under Assumptions 1–3 and* $\mu_1 > 0$*, there exists a unique interior* $\kappa^* \in (0,\infty)$ *maximizing* $W(\kappa)$*. The market equilibrium* $\kappa^{mkt} > \kappa^*$ *because firms do not internalize:* (i) *the congestion externality on workers' job-finding rates;* (ii) *the irreversibility of permanent exit to D;* (iii) *the wage depression imposed on non-routine incumbents and the consequent slowing of retraining via competing hazards on U.*

Assumption 4 (Irreversibility premium)**.** $\chi^U > 0$ *and* $\chi^D > \chi^U$*. Permanent exit is costlier per unit time than temporary unemployment.*

Assumption 5 (Sufficient concavity)**.** $\chi^D > \bar{\chi}^D$*, where* $\bar{\chi}^D$ *is the explicit threshold in Appendix B, eq. (B.4). This ensures transition cost convexity dominates output gain concavity.*

Proof sketch. Three steps (full argument in Appendix B).

(1) $V_Y(\kappa) \equiv \int e^{-\beta t} Y \, dt$ is concave: logistic output gains saturate ($\partial^2 a_t / \partial \kappa^2 < 0$ for $a_t > 1/2$).

(2) $V_U(\kappa) \equiv \int e^{-\beta t} U \, dt$ is convex: $\Phi(t;\kappa)$ is convex in κ for all t, and $\mu_1 > 0$ ensures hazard rates do not fall in κ.

(3) $V_D(\kappa) \equiv \int e^{-\beta t} D \, dt$ is convex: both μ and U increase in κ during the transition, making $\mu \cdot U$ superlinear in κ. The competing hazards structure amplifies this: higher μ depletes U, slowing $R(t)$ and further depressing w^N, reinforcing μ in a self-amplifying loop.

$W = V_Y - \chi^U V_U - \chi^D V_D$ is strictly concave under Assumption 3. $W'(0) > 0$ (foregone output). $W'(\kappa) < 0$ for large κ: logistic saturation implies $V_Y'(\kappa) \to 0$, while peak displacement $\Phi_{\max} = \bar{\iota}\kappa/4$ scaling linearly in κ with bounded exit hazard $\mu \leq \bar{\mu}$ implies the marginal unemployment cost $\chi^U V_U'(\kappa)$ remains bounded away from zero for all large κ (Lemma B.1 in Appendix B). The $\chi^D V_D'$ term reinforces this. Unique κ^* by strict concavity and IVT. □

Corollary 1 (Comparative statics of κ^*)**.**

Change	κ^* *moves*	*Mechanism*
$\bar{\iota} \uparrow$	$\kappa^* \downarrow$	*More workers at risk — go slower*
$\rho(Z) \uparrow$	$\kappa^* \uparrow$	*Queue clears faster — can absorb more*
$\chi^D \uparrow$	$\kappa^* \downarrow$	*Permanent exit costlier — go slower*
$\beta \uparrow$	$\kappa^* \uparrow$	*Impatient planner front-loads gains*
$\sigma \downarrow$	$\kappa^* \downarrow$	*Wage depression steeper — go slower*

Prediction 1 (Unemployment hump). Peak unemployment $U_{\max,c}$ is increasing in κ_c and $\bar{\imath}_c$, decreasing in $\rho(Z_c)$, with interaction $\kappa_c \times \bar{\imath}_c$ capturing exposure amplification. High-$\bar{\imath}$ countries suffer disproportionately under fast adoption.

Prediction 2 (LFP path). The welfare-relevant object is not the terminal LFP level but the *path*: how steeply participation falls and how early. Countries with high $\kappa_c \times \bar{\imath}_c$ should exhibit earlier and steeper LFP declines, particularly among older and less-educated workers (higher implicit χ^D), conditional on $\rho(Z_c)$.

Prediction 3 (Non-routine wage compression). Non-routine wages should compress earlier and more steeply in high-κ countries, particularly at the bottom of the wage distribution where retrained workers enter. This prediction distinguishes the mechanism from skill-biased technical change, which raises non-routine wages.

Prediction 4 (Labor share path). The discounted labor income share should decline faster and reach its trough earlier in high-κ countries, even controlling for the long-run automation level $\bar{\imath}$. The long-run labor share difference across κ is second-order; the transition-path difference is first-order.

Prediction 5 (Policy complementarity). Countries that increase $\rho(Z)$ before or during fast adoption should exhibit smaller LFP declines and earlier wage recovery than countries that adopt fast without institutional investment. This is the cross-country analogue of $\partial\kappa^*/\partial\rho(Z) > 0$.

7 Extension: Endogenous Task Creation and Two-Speed Dynamics

The baseline model abstracts from new task creation, holding the non-routine task measure fixed at $(1 - \bar{\imath})$. This section introduces a second diffusion process governing the speed of task creation and shows how it interacts with adoption speed $\kappa_A \equiv \kappa$ to jointly determine wages, permanent exit, and optimal policy. The baseline is the special case $\kappa_T = 0$.

7.1 Task Creation Process

Let $B(t)$ denote the cumulative measure of new non-routine tasks created since $t = 0$, following its own logistic diffusion:

$$\dot{b}_t = \kappa_T \cdot b_t(1 - b_t), \qquad b_0 \in (0, 1) \tag{25}$$

with $B(t) = \bar{b} \cdot b_t$ and $\bar{b} > 0$ the long-run measure of new tasks created. $\kappa_T > 0$ is the *task-creation speed* — the rate at which firms reorganize, create complementary roles, and expand the frontier of non-routine work. The non-routine task measure becomes:

$$M_N(t) = (1 - \bar{\imath}) + B(t) \tag{26}$$

expanding over time as new tasks are created. The baseline model corresponds to $\kappa_T = 0$ (or equivalently $\bar{b} = 0$): no task creation, fixed task frontier.

7.2 Modified Wage Equation

With expanding task measure $M_N(t)$, the symmetric task assignment becomes $N(i,t) = L^N(t)/M_N(t)$ for $i \in (\bar{\imath}, \bar{\imath} + M_N(t))$. The market-clearing wage is now determined by the ratio of task demand to labor supply:

$$w^N(t) = w_0^N \cdot \left[\frac{M_N(t)}{L^N(t)}\right]^{1/\sigma} \tag{27}$$

replacing (15). Three forces now govern $w^N(t)$:

- **Labor supply pressure** ($L^N \uparrow$): retraining adds workers, depressing w^N as before.
- **Task expansion** ($M_N \uparrow$): new tasks create demand for non-routine labor, pushing w^N upward.
- **Net effect**: w^N falls iff labor supply grows faster than task creation — iff $\dot{L}^N/L^N > \dot{M}_N/M_N$.

The wage depression mechanism of the baseline model therefore operates when κ_A is large relative to κ_T: fast adoption floods the labor market before the task frontier can absorb it.

7.3 Two-Speed Proposition

Proposition 5 (Two-Speed Interaction)**.** *In the extended model with task creation* (25)–(27)*:*

(i) *$D(\infty)$ is increasing in κ_A and decreasing in κ_T, holding the other fixed.*

(ii) *The optimal adoption speed κ_A^* is increasing in κ_T: economies with faster task creation can sustain faster AI adoption without welfare losses.*

(iii) *The welfare wedge $\kappa_A^{mkt} - \kappa_A^*$ is decreasing in κ_T: the case for slowing adoption is weakest in high task-creation economies.*

(iv) *As $\kappa_T \to \infty$, task creation occurs instantaneously up to $\bar{b}$, so $B(t) \to \bar{b}$ for all $t > 0$ and $M_N(t) \to (1-\bar{\imath})+\bar{b}$. Wage depression is eliminated* if and only if *$\bar{b}$ is sufficiently large: specifically, $(1-\bar{\imath})+\bar{b} \geq L_\infty^N$ is required for $w^N(t) \geq w_0^N$ to hold. Fast task creation alone is insufficient — it is the* scale *$\bar{b}$, not merely the speed κ_T, that determines whether the wage channel is eliminated. As $\bar{b} \to \infty$ (unbounded task frontier): $w^N(t) \to w_0^N$, the wage depression channel $\mu_2 < 0$ is inoperative, $D(\infty) \to D_{\min} > 0$ (from congestion alone), and the model converges to the Acemoglu–Restrepo benchmark where the task frontier expands without limit to absorb displaced workers.*

Proof sketch. Parts (i)–(ii). Higher κ_T raises $\dot{B}(t)$ for all t, expanding $M_N(t)$ faster. From (27), this attenuates wage depression: $\partial w^N / \partial \kappa_T > 0$, which reduces μ via $\mu_2 < 0$, reducing $D(\infty)$ via the integral $\int \mu \cdot U \, ds$. In the welfare problem (23), the marginal social cost of κ_A falls as κ_T rises (wage depression less severe), shifting κ_A^* upward. *Part (iii)* follows directly. *Part (iv):* as $\kappa_T \to \infty$, $B(t) \to \bar{b}$ instantaneously, so $M_N(t)$ dominates $L^N(t)$ and $[M_N/L^N]^{1/\sigma} \to \infty/\infty$. At the balanced path where $B(t)/L^N(t) \to$ const, $w^N(t) \to w_0^N$ recovers. With $\mu_2 < 0$ inoperative, only the congestion channel ($\mu_1 > 0$) generates discouragement, yielding $D_{\min} = \int \mu(q, w_0^N) \cdot U \, ds > 0$. □

Remark 5 (Nesting Acemoglu–Restrepo)**.** *Part (iv) provides a clean nesting of the AR framework. In AR, reinstatement is sufficiently vigorous that new task creation absorbs displaced workers at the long-run equilibrium wage. This corresponds to $\kappa_T \to \infty$ in our model: the task frontier expands without limit, eliminating wage depression. Our baseline ($\kappa_T = 0$) and the AR benchmark ($\kappa_T \to \infty$) are the two polar cases; the empirically relevant region is intermediate, where κ_T is positive but finite and the relative magnitudes of κ_A and κ_T determine transition costs. The key insight our model adds to AR is that even at the AR long-run equilibrium, the* path *to that equilibrium generates permanent welfare losses if $\kappa_A \gg \kappa_T$ during the transition.*

Corollary 2 (Extended comparative statics of κ_A^*)**.** *Corollary 1 gains one row:*

Change	*κ_A^* moves*	*Mechanism*
$\kappa_T \uparrow$	$\kappa_A^* \uparrow$	*Task creation absorbs workers — wage channel weakens*

7.4 Empirical Implications of the Two-Speed Model

The extension generates a fourth cross-country prediction:

Prediction 4 (Task-creation moderation). The negative effect of $\kappa_A \times \bar{\imath}$ on LFP (Prediction 2) and the wage compression effect (Prediction 3) should both be smaller in magnitude in countries with faster task creation κ_T. Formally:

$$\Delta \text{LFP}_c = \beta_0 - \beta_1(\kappa_{A,c} \times \bar{\imath}_c) + \beta_2 \rho(Z_c) + \beta_3 \kappa_{T,c} + \beta_4(\kappa_{A,c} \times \bar{\imath}_c \times \kappa_{T,c}) + \varepsilon_c$$

Expected: $\beta_4 > 0$ (task creation moderates the displacement-exit relationship). $\kappa_{T,c}$ is proxied by growth in new non-routine job titles or AI-complementary skill clusters from online job postings (Burning Glass/Lightcast, LinkedIn/Revelio), capturing the rate at which new roles emerge in the non-routine sector.

Option B extension (future work). A natural extension is to endogenize κ_T by making task creation respond to automation rents: $\dot{B} = \eta A(t) - \delta B(t)$, where faster adoption raises profits and accelerates complementary task creation. This creates a second feedback loop through which κ_A partially offsets its own displacement effect, making the net welfare impact of adoption speed an empirical question about relative magnitudes. We leave this to future work.

8 Conclusion

This paper has studied a dimension of AI adoption that steady-state analyses cannot address: not how much gets automated, but how fast. Holding fixed the long-run automation envelope $\bar{\imath}$—the Acemoglu–Restrepo destination—we have shown that adoption speed κ is the sole determinant of transition welfare. Two economies that arrive at the same long-run equilibrium can experience radically different labor market outcomes along the way, and the one that arrives faster is not necessarily better off. It may be permanently worse off.

The mechanism is the shadow price of time in the retraining queue. Fast adoption compresses the displacement window: the same total mass $\bar{\imath}$ of displaced workers arrives in a shorter period, overwhelming retraining capacity $\rho(Z)$. Workers who rationally update on lengthening queues and depressed non-routine wages exit the labor force permanently rather than waiting for a slot that may not arrive. This exit is absorbing. The welfare cost of fast adoption is therefore not measured by where labor force participation ends up, but by how steeply it falls

and how long it stays depressed—and by the sustained compression of the labor share throughout the transition window.

Four results characterize this mechanism. Peak unemployment is strictly increasing in κ: concentration of displacement into a shorter window raises the peak without changing total displacement $\int \Phi \, dt = \bar{\imath}$. Permanent labor force exit is larger under fast adoption even when terminal participation levels are similar across paths, because the discounted labor share is unambiguously lower throughout the transition. Non-routine wages follow a crossing pattern—lower during the transition window under fast adoption, then partially recovering as the discouraged stock depletes the retrained workforce—so that faster adoption can raise long-run wages for non-routine workers while permanently reducing overall participation, an outcome that wage-level statistics alone would mischaracterize as beneficial. And social welfare is strictly concave in κ, maximized at an interior optimum $\kappa^* < \kappa^{\text{mkt}}$, because individual firms do not internalize the congestion they impose on other workers' retraining queues, the irreversibility of permanent exit, or the wage depression borne by non-routine incumbents. These results pertain to the transition phase of AI adoption, when displacement and retraining dynamics dominate. As diffusion progresses and productivity gains materialize, these effects may attenuate or reverse.

The model identifies two distinct levers, each operating through a different channel. Retraining capacity $\rho(Z)$—the throughput rate of the economy's reallocation pipeline—is shifted by active labor market programs, retraining subsidies, credential reform, and mobility support. The timing of this investment matters: the comparative statics of the congestion index $q(t) = \Phi(t;\kappa)/(\rho(Z) \cdot U(t))$ show that institutions built *before* the displacement peak are far more effective than those built in response to it, because the discouragement hazard μ is highest precisely when congestion peaks.

A complementary instrument operates at the firm level. If firms sponsor retraining for a share $\theta \in [0,1]$ of their displaced routine workers—through internal redeployment, apprenticeship schemes, or subsidized upskilling—those workers bypass the public queue entirely, reducing effective inflow to $(1-\theta)\Phi(t;\kappa)$ and attenuating congestion without requiring expansion of institutional capacity $\rho(Z)$. Unlike public ALMP investment, firm-sponsored retraining is congestion-free by construction: it does not compete for slots in the shared pipeline. This suggests that policies subsidizing firm-level retraining—retention credits, co-investment mandates, or tax incentives tied to displacement events—are complements to, not substitutes for, public retraining capacity.

The two instruments address different margins: $\rho(Z)$ governs the throughput of the reallocation pipeline for workers already displaced; θ reduces the displacement load on that pipeline before it arrives. Adoption speed κ itself is not a physical constant but the aggregate outcome of decentralized firm decisions; within limits it can be shaped by AI liability frameworks, sector-specific deployment standards, and the timing of public procurement. The two instruments are not interchangeable: $\rho(Z)$ addresses the reallocation margin and κ addresses the displacement margin, and the welfare optimum requires both to be set simultaneously.

Several abstractions deserve scrutiny. The model treats the automation threshold $\bar{\imath}$ as exogenous and fixed; in reality it evolves with AI capabilities, factor prices, and regulatory choices. The reallocation hazard $\rho(Z)$ is treated as institutionally determined and independent of wages, an assumption that simplifies the analysis but suppresses potential crowding-in effects of high non-routine wages on retraining investment. The discouragement hazard μ is microfounded through an optimal stopping problem, but the aggregate exit dynamics depend on the distribution of individual thresholds in ways that the representative-worker formulation cannot fully capture.

The empirical agenda follows directly from the cross-country predictions of Section 6. The central challenge is measuring κ at the country or sector level—a logistic diffusion curve fitted to harmonized firm-level AI adoption data—and isolating its causal effect on LFP paths using pre-period routine task exposure as a Bartik-style instrument. The five predictions are sharp enough to be falsified: Prediction 2, in particular, requires not just a negative level effect on LFP but an earlier and steeper *path*, a richer test than any static cross-section can provide. We leave the empirical implementation to companion work.

The conclusion that emerges is simple and, we believe, underappreciated. The policy debate over AI and labor markets has asked the wrong question. The right question is not how much will be automated but whether our institutions can absorb automation at the speed the market has chosen. Speed is the variable. Capacity is the constraint. And the cost of getting the sequencing wrong is not temporary dislocation but permanent exit—workers who lose the train and never find another.

References

Daron Acemoglu. Directed technical change. *Review of Economic Studies*, 69(4): 781–809, 2002.

Daron Acemoglu and Pascual Restrepo. The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review*, 108(6):1488–1542, 2018.

Daron Acemoglu and Pascual Restrepo. Tasks, automation, and the rise in US wage inequality. *Econometrica*, 90(5):1973–2016, 2022.

David H. Autor. Work of the past, work of the future. *AEA Papers and Proceedings*, 109:1–32, 2019.

David H. Autor and David Dorn. The growth of low-skill service jobs and the polarization of the US labor market. *American Economic Review*, 103(5):1553–1597, 2013.

David H. Autor, David Dorn, and Gordon H. Hanson. The china syndrome: Local labor market effects of import competition in the United States. *American Economic Review*, 103(6):2121–2168, 2013.

David H. Autor, Caroline Chin, Anna M. Salomons, and Bryan Seegmiller. New frontiers: The origins and content of new work, 1940–2018. *Quarterly Journal of Economics*, 139(3):1399–1465, 2024.

Hermann Brunner. *Collocation Methods for Volterra Integral and Related Functional Equations*. Cambridge University Press, 2004.

Erik Brynjolfsson, Daniel Rock, and Chad Syverson. Artificial intelligence and the modern productivity paradox: A clash of expectations and statistics. *NBER Working Paper*, (24001), 2017.

Lorenzo Caliendo, Maximiliano Dvorkin, and Fernando Parro. Trade and labor market dynamics: General equilibrium analysis of the China trade shock. *Econometrica*, 87(3):741–835, 2019.

Diego Comin and Bart Hobijn. An exploration of technology diffusion. *American Economic Review*, 100(5):2031–2059, 2010.

Rafael Dix-Carneiro. Trade liberalization and labor market dynamics. *Econometrica*, 82(3):825–885, 2014.

Michael W. L. Elsby, Bart Hobijn, and Aysegul Sahin. The decline of the US labor share. *Brookings Papers on Economic Activity*, 44(2):1–63, 2013.

Claudia Goldin and Lawrence F. Katz. *The Race between Education and Technology*. Harvard University Press, 2008.

Jonathan Heathcote, Kjetil Storesletten, and Giovanni L. Violante. Optimal tax progressivity: An analytical framework. *Quarterly Journal of Economics*, 132 (4):1693–1754, 2017.

Nir Jaimovich and Henry E. Siu. The trend is the cycle: Job polarization and jobless recoveries. *NBER Working Paper*, (18334), 2012.

Richard K. Miller and Anthony N. Michel. *Ordinary Differential Equations*. Academic Press, New York, 1982.

Jan Tinbergen. Substitution of graduate by other labour. *Kyklos*, 28(2):217–226, 1975.

Appendix

A CES Production Structure and Wage Derivation

Output is produced by combining tasks $i \in [0,1]$ via CES aggregation:

$$Y_t = \left[\int_0^{A(t)} (A_K k(i))^{(\sigma-1)/\sigma} di + \int_{A(t)}^{\bar{\imath}} (A_R R(i))^{(\sigma-1)/\sigma} di + \int_{\bar{\imath}}^{1} (A_N N(i))^{(\sigma-1)/\sigma} di \right]^{\sigma/(\sigma-1)} \tag{A.1}$$

where $A(t) = \bar{\imath} \cdot a_t$ is the automated task share, A_K, A_R, A_N are productivity parameters, and r is the rental rate of AI capital. $Y(t)$ is increasing in $A(t)$ when $A_K/r > A_R/w^R$ (AI capital cheaper than routine labor), and increasing in $L^N(t)$ via the third integral.

Under symmetric task assignment within the non-routine region—$N(i,t) = L^N(t)/(1-\bar{\imath})$ for $i \in (\bar{\imath}, 1]$—the marginal product of non-routine labor is:

$$w^N(t) = \frac{\partial Y}{\partial L^N} = A_N^{1-1/\sigma} \cdot \left(\frac{Y_t}{L^N(t)} \right)^{1/\sigma} \cdot (1-\bar{\imath})^{1/\sigma} \tag{A.2}$$

Evaluating at $t = 0$ where $L^N(0) = 1 - \bar{\imath}$ gives w_0^N. Setting $L^{N,d}(w^N(t)) = L^N(t)$ in market clearing and using $Y_t/Y_0 \approx 1$ during early transition (AI capital share initially small) yields the reduced-form equilibrium wage (15). The approximation is exact when β is large.

The labor demand elasticity is $-\sigma$: a 1% increase in w^N reduces firm demand for non-routine labor by $\sigma > 1$ percent. Conversely, a 1% increase in labor supply (from retraining) depresses w^N by $1/\sigma < 1$ percent. The non-routine labor market absorbs all retrained workers through this price mechanism.

B Welfare Concavity

We prove $W(\kappa) = V_Y - \chi^U V_U - \chi^D V_D$ is strictly concave without invoking unbounded μ or the limit $\kappa \to \infty$.

Step 1: $V_Y(\kappa)$ is concave. $Y(t;\kappa)$ is increasing in $a_t(\kappa)$. For fixed t, $\partial a_t/\partial \kappa > 0$ (faster adoption $\Rightarrow$ more automation by t). $\partial^2 a_t/\partial \kappa^2 < 0$ for $a_t > 1/2$ (the logistic saturates post-inflection). Integrating with $e^{-\beta t}$, $V_Y''(\kappa) < 0$ when β is not too

small.

Step 2: $V_U(\kappa)$ is convex. From the Duhamel representation, $V_U(\kappa)$ is a positive linear functional of $\Phi(\cdot;\kappa)$ with a positive kernel. Convexity of V_U in κ follows if $\partial^2\Phi/\partial\kappa^2 \geq 0$ on a set of positive measure. For $t \leq t^*(\kappa)$ (the pre-inflection phase, $a_t \leq 1/2$), the logistic implies $\partial[\kappa a_t(1-a_t)]/\partial\kappa > 0$ and $\partial^2[\kappa a_t(1-a_t)]/\partial\kappa^2 > 0$: displacement inflow is strictly convex in κ during the build-up phase. The post-inflection phase contributes a negative term, but since the pre-inflection window has positive discounted measure and $e^{-\beta t}$ weights it most heavily, $V_U''(\kappa) > 0$ holds for β not too large.

Step 3: $V_D(\kappa)$ is convex. $D(t;\kappa) = \int_0^t \mu(q, w^N) \cdot U \, ds$. The integrand $\mu \cdot U$ involves two quantities both increasing in κ during the transition: U (by Step 2) and μ (via $\mu_1 > 0$ and $\mu_2 < 0$). The competing hazards structure amplifies convexity: higher μ depletes U, which slows $R(t)$, further depressing w^N and raising μ again. The cross-derivative $\partial^2(\mu U)/\partial\kappa^2 > 0$ (superlinear congestion). Hence $V_D''(\kappa) > 0$.

Conclusion and threshold. $W''(\kappa) = V_Y'' - \chi^U V_U'' - \chi^D V_D'' < 0$ provided:

$$\chi^D > \bar{\chi}^D \equiv \frac{|V_Y''(\kappa^*)| - \chi^U V_U''(\kappa^*)}{V_D''(\kappa^*)} \tag{B.4}$$

This is Assumption 3. The threshold $\bar{\chi}^D$ is finite and positive ($V_D'' > 0$ by Step 3).

$W'(0) > 0$ follows from foregone output: at $\kappa = 0$ no automation occurs and output gains are entirely unrealized, so the marginal benefit of faster adoption is positive.

$W'(\kappa) < 0$ for sufficiently large κ is established by the following lemma, which avoids any claim that $V_U(\kappa) \to \infty$ (a potentially false statement given bounded hazards and bounded displaced mass $\leq \bar{\imath}$).

Lemma 1 (Marginal unemployment cost is bounded below)**.** *There exists $\kappa_0 < \infty$ and $c > 0$ such that for all $\kappa \geq \kappa_0$,*

$$V_U'(\kappa) \;\equiv\; \frac{d}{d\kappa}\int_0^\infty e^{-\beta t} U(t;\kappa)\, dt \;\geq\; c \;>\; 0.$$

Proof. The proof avoids differentiating through the endogenous hazard system by working entirely with a linear comparison solution.

Step 1: ODE comparison. Under the logit hazard (20), $\mu(t) \leq \bar{\mu}$ for all t. Let $\widetilde{U}$ solve the linear comparison ODE

$$\dot{\widetilde{U}} = \Phi(t;\kappa) - (\rho + \bar{\mu})\,\widetilde{U}, \qquad \widetilde{U}(0) = 0.$$

Since $\mu(t) \leq \bar{\mu}$ implies U drains no faster than $\widetilde{U}$, the standard ODE comparison theorem gives $U(t;\kappa) \geq \widetilde{U}(t;\kappa) \geq 0$ for all t.

Step 2: Fubini identity. The variation-of-parameters solution is $\widetilde{U}(t) = \int_0^t e^{-(\rho+\bar{\mu})(t-s)}\Phi(s;\kappa)\,ds$. Integrating with discount factor and switching order of integration (Fubini, justified since the integrand is non-negative):

$$\int_0^\infty e^{-\beta t}\widetilde{U}(t;\kappa)\,dt = \frac{1}{\beta+\rho+\bar{\mu}}\int_0^\infty e^{-\beta s}\Phi(s;\kappa)\,ds. \tag{28}$$

Hence $V_U(\kappa) \geq \frac{1}{\beta+\rho+\bar{\mu}}\int_0^\infty e^{-\beta s}\Phi(s;\kappa)\,ds \equiv \widetilde{V}_U(\kappa)$.

Step 3: Derivative of the inflow integral. It suffices to show $\widetilde{V}_U'(\kappa) \geq c_0 > 0$ for all large κ, since $V_U'(\kappa) \geq \widetilde{V}_U'(\kappa)$ follows from Step 2. Differentiating under the integral by dominated convergence: since $a_s \in [0,1]$ and $\partial_\kappa a_s = s\,a_s(1-a_s) \leq s/4$, the derivative $\partial_\kappa[\kappa a_s(1-a_s)] = a_s(1-a_s) + \kappa(1-2a_s)\partial_\kappa a_s$ is bounded by an integrable envelope $C(1+s)$, and $e^{-\beta s}(1+s)$ is integrable for any $\beta > 0$:

$$\widetilde{V}_U'(\kappa) = \bar{\imath}\int_0^\infty e^{-\beta s}\,g(s,\kappa)\,ds.$$

At $s = 0$, $a_0 \in (0,1)$ is fixed, so $g(0,\kappa) \equiv \partial_\kappa[\kappa a_0(1-a_0)] = a_0(1-a_0) > 0$ for all κ. By continuity of $g(s,\kappa) \equiv \partial_\kappa[\kappa\,a_s(1-a_s)]$ jointly in (s,κ) and the fact that $g(0,\kappa) = a_0(1-a_0) > 0$ for all κ, there exist $\delta > 0$ and $\kappa_0 < \infty$ such that for all $\kappa \geq \kappa_0$ and all $s \in [0,\delta]$:

$$g(s,\kappa) \;\equiv\; \frac{\partial}{\partial\kappa}\big[\kappa\,a_s(1-a_s)\big] \;\geq\; \tfrac{1}{2}\,a_0(1-a_0) \;>\; 0.$$

Therefore, restricting to $[0,\delta]$ and bounding $e^{-\beta s} \geq e^{-\beta\delta} > 0$:

$$\widetilde{V}_U'(\kappa) \;\geq\; \bar{\imath}\cdot e^{-\beta\delta}\cdot\tfrac{1}{2}\,a_0(1-a_0)\cdot\delta \;\equiv\; c_0 \;>\; 0.$$

Setting $c \equiv c_0/(\beta+\rho+\bar{\mu}) > 0$ completes the proof: $V_U'(\kappa) \geq c$ for all $\kappa \geq \kappa_0$. □

Remark 6. *The economic intuition is that at early times ($s \in [0,\delta]$), faster adoption unambiguously raises the displacement inflow — $\partial_\kappa[\kappa a_s(1-a_s)] > 0$ because $a_s \approx a_0$ is not yet near its ceiling. This early-time contribution to the*

discounted unemployment cost is bounded below by a positive constant regardless of what happens at the inflection or later. The comparison solution $\widetilde{U}$ makes this precise without requiring monotonicity of U in κ globally.

Completing the high-κ argument. Since $V_Y(\kappa) \leq \bar{V}_Y \equiv Y(\bar{\imath})/\beta < \infty$, logistic saturation implies $V_Y'(\kappa) \to 0$ as $\kappa \to \infty$. By Lemma 1, $\chi^U V_U'(\kappa) \geq \chi^U c > 0$ for all large κ. Hence $W'(\kappa) = V_Y'(\kappa) - \chi^U V_U'(\kappa) - \chi^D V_D'(\kappa) < 0$ for sufficiently large κ, since the marginal output gain vanishes while the marginal unemployment cost remains bounded away from zero. The $\chi^D V_D'(\kappa)$ term is non-negative (Step 3) and only reinforces this. Unique κ^* follows by strict concavity and the IVT applied to W'.

C Constant-Hazard Approximation for $D(\infty)$

Under the auxiliary assumption that ρ and μ are constant at their peak-displacement values, and using the linear specification (22), the implied steady-state exit share is:

$$D(\infty) \approx \frac{\bar{\mu}}{\rho(Z) + \bar{\mu}} \cdot \bar{\imath} \tag{C.1}$$

where $\bar{\mu} = \mu_0 + \alpha_1 q^* - \alpha_2 w^N(t^*)$ is the discouragement hazard evaluated at peak congestion $(q^*, w^N(t^*))$. This is a calibration benchmark only—not an exact result of the general model. Given observed $D(\infty)$, $\bar{\imath}$, $\rho(Z)$, q^*, and $w^N(t^*)$, one can identify μ_0, α_1, and α_2 jointly with cross-country variation in κ and σ as described in Appendix D.

D Calibration Targets

Parameter	Range	Source / Identification
$\bar{\imath}$	0.30–0.45	OECD PIAAC routine task share; Autor–Dorn classification
κ	0.3–0.6	Logistic fits to AI adoption surveys (McKinsey, Stanford AI Index)
$\rho(Z)$	0.2–0.5/yr	ALMP spending, EPL index, job-finding rates (OECD)
σ	1.5–4	Task substitution elasticities (micro estimates); baseline $\sigma = 1.5$ in Fig
β	0.03–0.05	Standard macro discount rate
$s_{K,0}$	0.03–0.07	AI capital income share at start of transition (EU KLEMS)
χ^U	normalize	Match UI replacement rates and job-search costs
χ^D	calibrate	Match observed κ^{mkt} vs κ^* wedge
μ_0	back out	Baseline exit rate: from $D(\infty)$ at $q = 0$, $w^N = w_0^N$ via (C.1)
α_1	back out	Queue sensitivity: from cross-country variation in $D(\infty)$ vs κ
α_2	back out	Wage sensitivity: from cross-country variation in $D(\infty)$ vs w^N compre

E Existence and Uniqueness of Equilibrium Paths

We establish that the coupled ODE–integral system has a unique solution path, resolving the apparent simultaneity between $w^N(t)$, $L^N(t)$, and $U(t)$.

Structure of the system. The wage $w^N(t)$ depends on $L^N(t) = (1 - \bar{\imath}) + \rho(Z) \int_0^t U(s)\, ds$, which is a functional of the *history* of U up to t, not of $U(t)$ simultaneously. The system is therefore a Volterra integro-differential equation in $U(t)$, not a simultaneous algebraic-differential system:

$$\frac{dU}{dt} = \Phi(t;\kappa) - \left[\rho(Z) + \mu\left(\frac{\Phi(t;\kappa)}{\rho(Z)\max\{U(t),\varepsilon\}},\ w_0^N \cdot \left(\frac{1-\bar{\imath}}{(1-\bar{\imath}) + \rho(Z)\int_0^t U(s)\, ds}\right)^{1/\sigma}\right)\right] U(t) \tag{E.1}$$

with $U(0) = 0$, $D(0) = 0$, and all other states initialized at their pre-automation values.

Assumption 6 (Lipschitz discouragement hazard)**.** *$\mu(q, w)$ is Lipschitz continuous in both arguments. Under the logit specification* (20)*, this holds globally since $|\Lambda'(x)| \le 1/4$ for all x, bounding the Lipschitz constant by $L = \bar{\mu}/4 \cdot \sup|\nabla(V_D - V_U)|$, which is finite given bounded Φ, $w^N \in [w_\infty^N, w_0^N]$, and $U \ge \varepsilon$. Under the linear approximation* (22)*, $L = \alpha_1\|\partial q/\partial U\|_\infty + \alpha_2\|\partial w^N/\partial L^N\|_\infty < \infty$ for $U \ge \varepsilon$ and $L^N \ge 1-\bar{\imath}$. In both cases μ is bounded: $\mu \in (0, \bar{\mu})$ under* (20) *and $\mu \ge 0$ under*

(22) *by the truncation* $\mu = \max(0,\, \mu_0 + \alpha_1 q - \alpha_2 w^N)$*, which ensures non-negativity for all* (q, w^N) *without imposing a restriction on* μ_0*.*

Proposition 6 (Existence and Uniqueness)**.** *Under Assumption 4, the system* (E.1) *has a unique continuous solution* $U(t)$ *on* $[0, \infty)$*. Given* $U(t)$*, the paths* $L^N(t)$*,* $w^N(t)$*,* $D(t)$*, and* $R(t)$ *are uniquely determined by* (5)–(8) *and* (15)*.*

Proof. Define the operator $\mathcal{T}[U](t)$ as the right-hand side of (E.1) integrated from 0, mapping continuous functions on $[0, T]$ to themselves. The Lipschitz condition on μ (Assumption 4), combined with Φ continuous and bounded and the $\max\{U, \varepsilon\}$ bound preventing singularity, ensures $\mathcal{T}$ is a contraction on $C([0, T], [0, \bar{\imath}])$ with the sup norm for sufficiently small T. Existence and uniqueness on $[0, T]$ follow from the Banach fixed-point theorem applied to Volterra operators [Brunner, 2004, Theorem 2.1]; see also Miller and Michel [1982] for the standard continuation theory. The solution extends to $[0, \infty)$ since $U(t)$ is bounded above by $\bar{\imath}$ at all t: from the accounting identity (4), $L^R(t) + U(t) + D(t) + R(t) = \bar{\imath}$ with all terms non-negative, so $U(t) \leq \bar{\imath}$ immediately. $U(t) \geq 0$ follows from $\Phi \geq 0$ and the absorbing nature of D. Once $U(t)$ is uniquely determined, all other state variables follow uniquely from their defining equations. □

Remark 7. *The key insight is that* $w^N(t)$ *depends on the* integral *of* U*, not on* $U(t)$ *directly, so there is no simultaneity at any instant* t*. The system is causal: past retraining flows determine today's wage, which determines today's exit rate, which determines tomorrow's retraining flows.*

Extension to the two-speed model (Section 6). When task creation is active ($\kappa_T > 0$), the system acquires the additional state variable $B(t)$ satisfying $\dot{B} = \kappa_T \cdot b_t(1 - b_t) \cdot \bar{b}$, which is a standalone logistic ODE with no dependence on the labor market block. It has a unique smooth solution $B(t) \in [0, \bar{b}]$ for all $t \geq 0$, determined entirely by κ_T, $\bar{b}$, and the initial condition $b_0 \in (0, 1)$. The wage equation becomes

$$w^N(t) = w_0^N \cdot \left[\frac{M_N(t)}{L^N(t)}\right]^{1/\sigma}, \qquad M_N(t) = (1 - \bar{\imath}) + B(t).$$

Since $B(t)$ is predetermined (smooth and bounded, independent of U), $M_N(t)$ enters the right-hand side of (E.1) as a known, bounded, Lipschitz function of time. The Volterra structure of the baseline system is therefore preserved: $w^N(t)$ still depends on the integral of U and on $B(t)$, both of which are bounded and Lipschitz.

The same contraction argument applies without modification, and existence and uniqueness of the extended ODE system follow immediately from the baseline proof. $D(t)$ and $R(t)$ remain uniquely determined by integration once $U(t)$ is known.

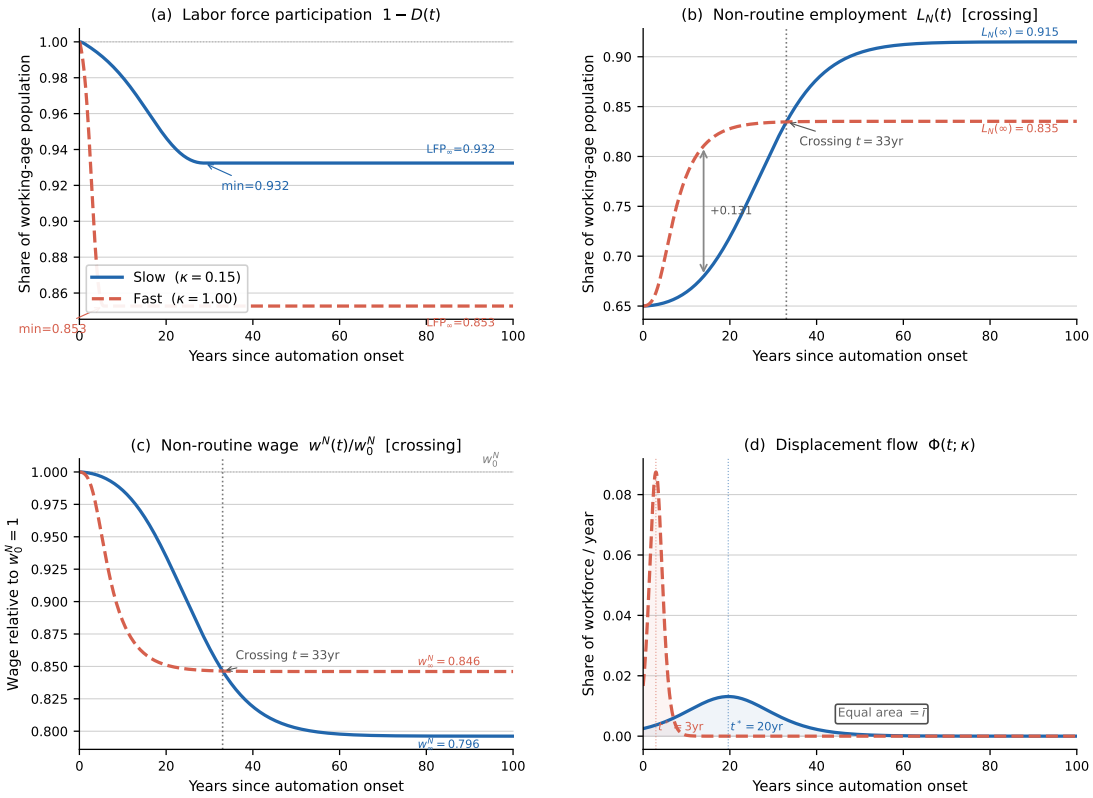


Figure 1: **Transition dynamics under slow vs. fast adoption.** Four panels show model paths for $\kappa = 0.15$ (slow, solid blue) and $\kappa = 1.00$ (fast, dashed red). Parameters: $\bar{\imath} = 0.35$, $\rho = 0.20$, $\sigma = 1.5$, $\mu_0 = 0.02$, $\alpha_1 = 0.12$, $\alpha_2 = 0.15$. The wage sensitivity $\alpha_2 = 0.15$ implies a 7.5 percentage point increase in cumulative exit probability per 10% wage shortfall over a five-year retraining spell, consistent with participation responses documented in Autor et al. [2013]. *Panel (a):* Labor force participation $1-D(t)$ — unambiguously lower under fast adoption throughout the path and permanently ($\text{LFP}_\infty = 0.853$ vs. 0.932, an 8 pp gap). *Panel (b):* Non-routine employment $L_N(t)$ — the crossing result. Fast adoption front-loads retraining inflow (peak gap +0.131 at $t \approx 14\,\text{yr}$) but higher permanent exit depletes the retrained stock; paths cross at $t \approx 33\,\text{yr}$ and $L_N(\infty)$ is 8 pp lower under fast adoption (0.835 vs. 0.915). *Panel (c):* Non-routine wage $w^N(t)/w_0^N$ — mirror image of (b). Wages fall more sharply under fast adoption during the transition window but end higher once the retrained stock is depleted ($w_\infty^N = 0.846$ vs. 0.796, crossing at $t \approx 33\,\text{yr}$). Faster adoption can permanently raise wages for surviving non-routine workers while permanently reducing overall participation. *Panel (d):* Displacement flow $\Phi(t;\kappa)$ — equal total displaced workers $\bar{\imath}$ under both paths, compressed into a 3-year window under fast adoption vs. 20-year window under slow adoption.

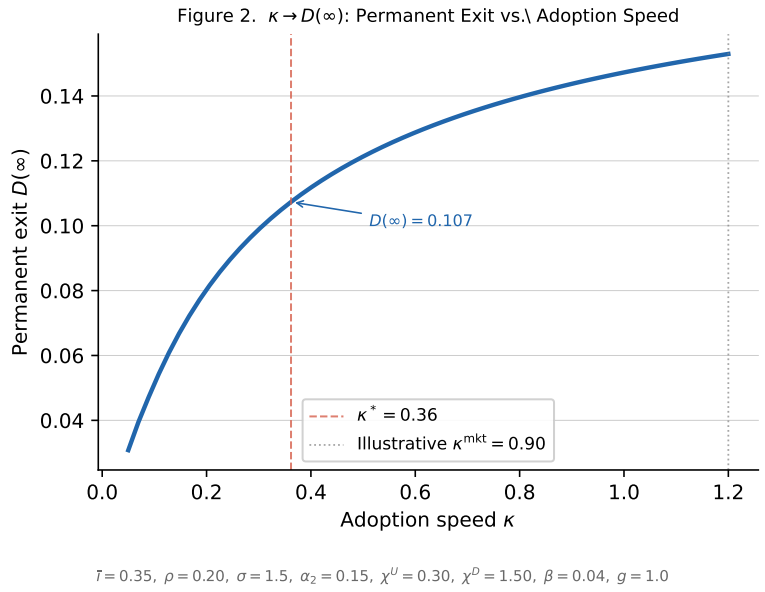


Figure 2: **Permanent exit $D(\infty)$ as a function of adoption speed κ.** The curve is concave and saturating: $D(\infty)$ rises steeply at low κ (congestion first becomes binding) and flattens at high κ (discouragement hazard near its ceiling). Under the unified calibration, $D(\infty)$ ranges from 3% at $\kappa = 0.05$ to 15% at $\kappa = 1.20$ — a fivefold increase. The dashed line marks $\kappa^* = 0.36$ and the dotted line marks an illustrative $\kappa^{\text{mkt}} = 0.90$; the gap is excess permanent exit from market over-adoption. All figures use the same parameter set: $\bar{\imath} = 0.35$, $\rho = 0.20$, $\sigma = 1.5$, $\mu_0 = 0.02$, $\alpha_1 = 0.12$, $\alpha_2 = 0.15$.

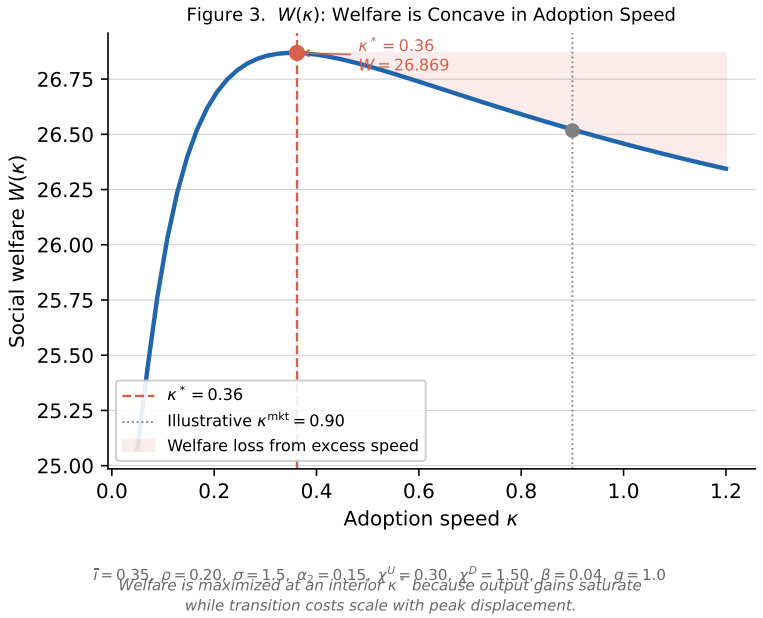


Figure 3: **Social welfare $W(\kappa)$ is concave in adoption speed.** Welfare is maximized at an interior $\kappa^* = 0.36$ because output gains saturate as the logistic diffusion completes, while transition costs scale with peak displacement $\Phi_{\max} = \bar{\imath}\kappa/4$. The shaded region is the welfare loss from adopting at the illustrative $\kappa^{\text{mkt}} = 0.90 > \kappa^*$. The planner's problem is not whether to automate but whether institutions can absorb automation at the chosen speed. Parameters unified across all figures: $\bar{\imath} = 0.35$, $\rho = 0.20$, $\sigma = 1.5$, $\alpha_2 = 0.15$, $\chi^U = 0.30$, $\chi^D = 1.50$, $\beta = 0.04$, $g = 1.0$.

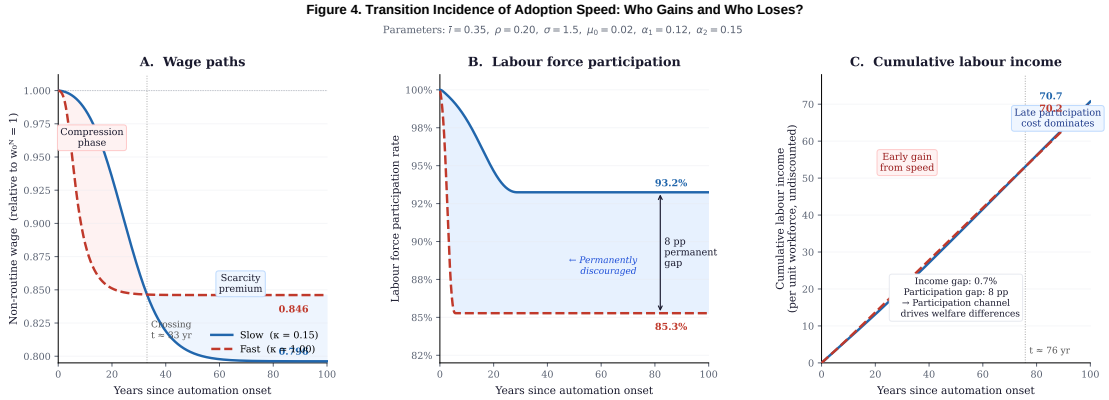


Figure 4: **Transition incidence of adoption speed** Three panels plot model paths for $\kappa = 0.15$ (slow, solid blue) and $\kappa = 1.00$ (fast, dashed red) under the unified parameter set $\bar{\iota} = 0.35$, $\rho = 0.20$, $\sigma = 1.5$, $\mu_0 = 0.02$, $\alpha_1 = 0.12$, $\alpha_2 = 0.15$. *Panel A:* Non-routine wages $w^N(t)/w_0^N$. During the compression phase ($t < 33$ yr, shaded red), fast adoption depresses wages as the retraining surge raises non-routine labour supply; the scarcity premium subsequently reverses the gap (shaded blue). Long-run wages are higher under fast adoption ($w_\infty^N = 0.846$ vs. 0.796) because permanent exit reduces the retrained stock. *Panel B:* Labour force participation $1 - D(t)$. The 8 pp permanent gap (shaded blue) reflects the discouraged-worker mechanism: workers who rationally exit during the congested retraining window do not return. *Panel C:* Cumulative undiscounted labour income $\int_0^t w^N(s) L_N(s)\, ds$ per unit workforce. Fast adoption leads until $t \approx 76$ yr, when the larger employed workforce under slow adoption overtakes. The terminal income gap is 0.7% — small relative to the 8 pp participation gap — indicating that the labour-income channel alone does not account for the welfare difference; the participation channel drives welfare differences across adoption speeds.