THE HETEROGENEOUS DIFFUSION OF AI: INDIVIDUALS, ORGANISATIONS, AND ADOPTION BARRIERS

Documents de travail GREDEG
GREDEG Working Papers Series

Hankui Wang
Jiachen Yi
Philipp Harting

GREDEG WP No. 2026-16
https://ideas.repec.org/s/gre/wpaper.html

The heterogeneous diffusion of AI: Individuals, organisations, and adoption barriers

Hankui Wang[a], Jiachen Yi[b], Philipp Harting[a]

[a] *Université Côte dAzur, CNRS, GREDEG, France,*
[b] *School of Economics, University of Bristol, United Kingdom,*

Abstract

Artificial Intelligence (AI) is adopted far faster by individuals than by organisations, yet most diffusion models treat both populations as homogeneous and independent. We develop an extended Bass model with heterogeneous archetypes, bidirectional cross-group spillovers, endogenous barrier decay, and productivity feedback. Base parameters are estimated from survey data on individual and organisational AI adoption across 45 countries (2020–2025); the extended model is calibrated and simulated over ten years. Two findings emerge: cross-group spillover from individual to organisational adoption is the dominant diffusion accelerator in the cross-country evidence, while within-group imitation is undetectable in the current two-year panel; and heterogeneity generates a 5.2-year spread across firm archetypes in time to 50% adoption, capturing most of the 6-year empirical firm-size gap. The calibrated model simulates overall organisational AI adoption reaching 50% by 2029 and 74% by 2033 under the benchmark calibration, translating into sustained productivity growth. Education and training delivers the largest adoption gain per unit of intervention intensity, while direct subsidies most effectively narrow the adoption gap between large and small firms.

Keywords: artificial intelligence, technology diffusion, Bass model

Email address: `hankui.wang@univ-cotedazur.fr` (Hankui Wang)

1. Introduction

Artificial Intelligence (AI) is a transformative general-purpose technology, yet its diffusion reveals a striking asymmetry between individual users and organisations. ChatGPT reached 100 million users within two months of release [1], the fastest consumer technology adoption on record. More than three years later, almost two-thirds of organisations have yet to scale AI to the enterprise, with the persistent gap reflecting integration challenges, strategic frictions, and uncertain returns [2]. Capturing these dynamics requires heterogeneity, cross-group spillovers, and productivity feedback.

This paper develops an extended Bass [3] diffusion model formulated as a system of coupled ordinary differential equations (ODEs). The model segments adopters into individual and organisational populations and incorporates four features: (1) heterogeneous adopter types, reflecting distinct incentives and frictions; (2) bidirectional cross-group influences, capturing spillovers between individuals and organisations; (3) a dynamic adoption barrier that erodes as diffusion advances; and (4) an endogenous productivity feedback loop, linking adoption to aggregate productivity.

The contribution of this paper is threefold. First, we jointly model these four features within a classical diffusion system, capturing the coupled feedback dynamics that characterise AI diffusion. Second, we estimate the base model's core parameters using microdata from the Pew Research Center, the US Census Bureau's Business Trends and Outlook Survey (BTOS), the OECD Information, Communication and Technology (ICT) Surveys, and Eurostat, spanning individual and organisational adoption across 45 countries (2020–2025). We then document two findings: (i) cross-country evidence is consistent with individual-to-organisational spillover as the dominant diffusion accelerator, while within-group imitation is undetectable in the current two-year panel; and (ii) adoption gaps across firm sizes and demographic groups reflect differences in *timing* rather than diffusion speed, closely tracking empirical gradients across both US and EU data. Third, we calibrate the full extended model from these estimates and auxiliary cross-sectional evidence, and simulate it forward. The calibrated model tracks observed adoption closely and shows that cross-group spillover, calibrated to the cross-country estimates, is the dominant accelerator of organisational diffusion in the simulation, with heterogeneity generating within-population timing dispersion that homogeneous models cannot represent. Policy simulations rank education and training as the most effective intervention overall, while direct subsidies are most effective at narrowing the firm-size adoption gap.

Section 2 reviews the literature. Section 3 presents the diffusion framework. Section 4 estimates parameters and documents heterogeneity patterns. Section 5 calibrates the model and evaluates policy interventions. Section 6 concludes.

2. Related literature

The technology diffusion literature provides the foundation for analysing AI adoption, originating with Griliches's (1957) study of hybrid corn, formalised in the Bass model [3], and extended to inter-firm [5, 6], intra-firm [7, 8], and multi-generation [9] settings. AI-specific applications adapt these frameworks to model adoption dynamics [10, 11, 12], while endogenous growth and productivity models offer complementary insights into long-term economic impacts [13, 14, 15, 16]. Agent-based models capture heterogeneous-agent dynamics through network-mediated interactions [17, 18]; our ODE framework trades that granularity for analytical tractability and empirical estimability while retaining heterogeneity through adopter archetypes. The bidirectional cross-group influences in our model reflect cross-side network effects, whereby two distinct user groups—here individuals and organisations—reinforce one another's adoption [19].

AI shares the defining features of a general-purpose technology [20], yet its diffusion differs from earlier general-purpose technologies (GPTs). Historical GPTs diffused gradually through organisational channels over decades

[21, 22]; the individual–organisational asymmetry documented above is unprecedented in scale and speed, and existing diffusion models offer no framework for analysing it.

Stylized Fact 1: *Individual AI adoption outpaces organisational deployment by a wide margin.*

Adoption varies systematically by firm size, sector, and national capabilities [12, 23]. Cross-country surveys report large-firm adoption rates several times higher than small-firm rates, concentrated within ICT and professional services; agriculture, construction, and accommodation lag, as do lower-income economies [24, 15, 23]. Common implementation challenges include integration with legacy systems, skill shortages, and data privacy concerns [24, 25].

Stylized Fact 2: *Larger firms and ICT-intensive sectors lead AI adoption; smaller firms lag.*

Two distinct micro-channels underpin spillover between segments. Employee familiarity with generative tools generates workforce expectations and demonstrated use cases that lower internal resistance to enterprise deployment [11, 10]; conversely, workplace exposure and the legitimising effect of enterprise adoption expand the population of consumer adopters [12]. Adoption trajectories across the two segments are therefore coupled rather than evolving independently.

Stylized Fact 3: *Countries with higher individual AI adoption also show higher organisational adoption.*

Adoption barriers vary across populations and over time. Cross-sectionally, they are concentrated on organisations—particularly smaller firms with limited absorptive capacity—and reflect skill gaps, integration costs, and organisational readiness [26, 27]. Dynamically, they erode as adoption itself advances, through three reinforcing channels: accumulated deployment experience lowers the cost of further adoption [28]; regulatory uncertainty diminishes as regulators issue guidance and legal precedents accumulate [29, 30];

and the supporting ecosystem of vendors and integration tools matures as the user base grows [20, 22]. Each channel strengthens with cumulative adoption, so the barrier facing remaining non-adopters falls endogenously as diffusion progresses.

Stylized Fact 4: *AI adoption barriers are higher for organisations and erode endogenously as diffusion progresses and complementary capabilities mature.*

Across these strands, AI's economic impacts are multifaceted. Productivity gains are positive but unevenly distributed, and do not uniformly translate to faster economic growth [14]. Research also highlights shifts in market structure and labour-market outcomes, including job displacement and evolving skill requirements [31, 32, 16]. Sector-specific studies document efficiency gains across healthcare, manufacturing, finance, and education [33, 25, 34, 35, 36, 37], though these gains appear only with substantial lags, as firms learn to use the technology and adapt their organisations [13, 14, 38].

Stylized Fact 5: *AI-related productivity gains are positive but uneven across sectors and materialise with time lags.*

Empirical evaluations from a range of country contexts identify three policy levers that systematically move adoption: cost-reducing subsidies for smaller firms [15, 12], training programmes that build absorptive capacity [32], and R&D incentives that strengthen the adoption-to-productivity link [37]. At the national level, digital readiness and aggregate R&D expenditure are among the strongest cross-country predictors of adoption [23].

Stylized Fact 6: *Targeted policy interventions can accelerate AI diffusion, particularly for lagging segments.*

Despite the breadth of evidence reviewed above, existing diffusion models treat heterogeneity, cross-group spillovers, and adoption barriers as independent features [7, 9, 6]. In practice, however, they reinforce each other: early-adopting firms lower barriers for later ones, while individual adoption spills over into organisational uptake. Capturing these feedback loops requires a joint framework, which we develop in Section 3 and estimate from cross-country survey data spanning 45 countries in Section 4.

3. Modelling framework

We model AI adoption as two coupled Bass-type populations—individuals and organisations—and then extend the base case to allow within-group heterogeneity, cross-group spillovers, and a productivity feedback loop.

3.1. A two-population Bass diffusion model

Reflecting the consumer-enterprise asymmetry documented in Stylized Fact 1, we segment the adopter population into individuals and organisations, modelling each within the Bass diffusion framework [3]. Let $F_{i,t}$ and $F_{o,t}$ denote the cumulative fractions that have adopted AI by time t. For individuals, the rate of adoption is the product of the adoption hazard—the probability that a non-adopter adopts at time t—and the remaining non-adopter fraction:

$$\frac{dF_{i,t}}{dt} = \underbrace{[p_i + q_i F_{i,t} + \alpha F_{o,t}]}_{\text{adoption hazard}}(1 - F_{i,t}), \tag{1}$$

where p_i is the *innovation* rate (autonomous adoption driven by external information), $q_i\, F_{i,t}$ is the *imitation* channel (within-group peer effects), and $\alpha\, F_{o,t}$ captures cross-group influence from organisational adoption (e.g., workplace exposure). The linear-additive hazard follows from a heterogeneous-cost threshold model derived in Appendix A.1. Organisational adoption follows an analogous structure:

$$\frac{dF_{o,t}}{dt} = [p_o + q_o F_{o,t} + \beta F_{i,t}]\,(1 - b_t)(1 - F_{o,t}), \tag{2}$$

with an analogous hazard—p_o, $q_o\, F_{o,t}$, and $\beta\, F_{i,t}$—but additionally modulated by $(1 - b_t)$, a time-varying barrier reflecting the frictions that slow organisational deployment relative to individuals. Consistent with Stylized Fact 4, the barrier erodes endogenously as adoption accumulates:

$$b_t = b_0\, e^{-k\,(F_{i,t}+F_{o,t})}. \tag{3}$$

where $b_0 \in [0, 1]$ is the initial barrier level and $k > 0$ is the erosion rate—the speed at which cumulative adoption reduces the remaining barrier. The exponential form embodies three properties: monotone decrease (every adoption erodes barriers), diminishing returns (early adoptions lower the barrier most), and strict positivity (residual frictions never fully vanish); alternative functional forms are discussed in Appendix A.1. Equal weighting of F_i and F_o is a parsimony assumption; Equation 6 in the extended model relaxes it with separate erosion coefficients for individual and organisational adoption.

The cross-group effects allow social influence to operate *between* populations rather than only within a single group as in the standard Bass model. The forward channel from individuals to organisations (βF_i) operates through workforce readiness [28, 8]. The reverse channel from organisations to individuals (αF_o) operates through workplace exposure [12].

In summary, the base model establishes the core dynamics: two populations adopting at different speeds, linked by bidirectional cross-group influence and subject to endogenously eroding barriers. Section 3.2 enriches this structure with within-group heterogeneity, cross-group influences, and a productivity feedback loop, while nesting the base case as a special case.

3.2. Heterogeneous extension with productivity feedback

The extended model generalises the base case along three dimensions while preserving it as a nested special case.

Individual-level heterogeneity. We partition individuals into G groups (e.g., early adopters and late adopters), indexed by g, each with population share

$s_{i,g}$, group-specific parameters $(p_{i,g}, q_{i,g})$, and adopted fraction $F_{i,g,t}$:

$$\frac{dF_{i,g,t}}{dt} = (p_{i,g} + q_{i,g}F_{i,t} + \alpha_g F_{o,t} + \lambda_{i,g}(P_t - 1))(1 - F_{i,g,t}), \tag{4}$$

where $F_{i,t} = \sum_{g=1}^{G} s_{i,g}F_{i,g,t}$ is aggregate individual adoption, $(1 - F_{i,g,t})$ is the non-adopter pool of group g. The hazard depends on aggregate signals—$F_{i,t}$, $F_{o,t}$, and P_t—that individuals observe economy-wide through peer use, news and products with embedded AI, and productivity gains, rather than on group- or archetype-specific quantities.

Organisational-level heterogeneity. To account for the firm-size heterogeneity documented in Stylized Fact 2, we partition organisations into M archetypes (e.g., SMEs, large enterprises), indexed by h, each with share $s_{o,h}$ and parameters $(p_{o,h}, q_{o,h})$:

$$\frac{dF_{o,h,t}}{dt} = (p_{o,h} + q_{o,h}F_{o,t} + \beta_h F_{i,t} + \lambda_{o,h}(P_t - 1))(1 - b_{h,t})(1 - F_{o,h,t}), \tag{5}$$

where $F_{o,t} = \sum_{h=1}^{M} s_{o,h}F_{o,h,t}$ is aggregate organisational adoption, $(1 - F_{o,h,t})$ is the non-adopter pool of archetype h, and $b_{h,t}$ is the archetype-specific barrier (defined below). The hazard structure parallels the individual case: aggregate signals $F_{i,t}$, $F_{o,t}$, and P_t enter through firms' awareness of economy-wide AI use and productivity gains.

Archetype-specific organisational barriers. Each archetype now has its own initial barrier $b_{0,h}$ and separate erosion coefficients for individual and organisational adoption:

$$b_{h,t} = b_{0,h} \exp\left[-(k_{o,h}F_{o,t} + k_{i,h}F_{i,t})\right], \tag{6}$$

where $b_{0,h}$ is the initial barrier for archetype h—analogous to the industry-specific fixed effects in Karshenas and Stoneman [6]—and $k_{o,h}$, $k_{i,h}$ govern how organisational and individual adoption erode it.

Saturating generalisation of cross-group influence. The base model's cross-group terms are linear in the other group's adoption. We generalise them to allow saturation, since cross-group influence plausibly saturates: early adopters send a strong viability signal, but beyond a critical mass the marginal informational gain diminishes. For the influence of individual adoption on organisations:

$$\beta_h F_{i,t} = \beta_{max,h} \left[1 - \exp(-\gamma_{i,h} F_{i,t})\right], \tag{7}$$

and for the influence of organisational adoption on individuals:

$$\alpha_g F_{o,t} = \alpha_{max,g} \left[1 - \exp(-\gamma_{o,g} F_{o,t})\right], \tag{8}$$

where $\beta_{max,h}$ and $\alpha_{max,g}$ are maximum influence strengths and γ governs the speed of saturation. The linear cross-group terms of the base model (Equations 1–2) are the $\gamma \to 0$ limit of Equations 7–8, with $\alpha = \alpha_{max,g}\gamma_{o,g}$ and $\beta = \beta_{max,h}\gamma_{i,h}$; under $G = M = 1$ and $\lambda = 0$ the extended model nests the base case.

Endogenous productivity feedback. Finally, to capture the lagged productivity effects of Stylized Fact 5, we formalise the productivity feedback loop. Let P_t denote a bounded, AI-specific index of productive capability [13], normalised so that $P_0 = 1$ represents the pre-AI baseline and bounded above by P_{max}. Its evolution is governed by

$$\frac{dP_t}{dt} = \varphi(s_i F_{i,t} + s_o F_{o,t})(P_{max} - P_t) - \delta P_t. \tag{9}$$

where the growth rate φ governs how fast adoption translates into capability gains; s_i and s_o weight the contributions of individual and organisational adoption; and $(P_{max} - P_t)$ yields diminishing returns near the ceiling P_{max}, consistent with Stylized Fact 5. The depreciation term δP_t captures knowledge obsolescence: deployed capabilities lose value without continued adoption, analogous to the depreciation of R&D capital [39]. The two terms balance at a steady state $P^* < P_{max}$ set by φ/δ: large φ pushes P^* near the ceiling, while large δ pulls it down. P_t then feeds back into adoption through the $\lambda_{i,g}(P_t - 1)$ and $\lambda_{o,h}(P_t - 1)$ terms (the gain over the pre-AI baseline $P_0 = 1$) in Equations 4 and 5, where $\lambda_{i,g}$ and $\lambda_{o,h}$ control how strongly each group and archetype responds to capability gains. This creates a self-reinforcing loop: adoption raises P_t, which raises the adoption hazard, accelerating further diffusion.

The extended model thus consists of a system of $G+M+1$ coupled equations: G equations for individual groups, M for organisational archetypes, and one for productivity growth. Together with the auxiliary equations for barriers and influences, this framework enables a richer exploration of AI diffusion under various heterogeneity assumptions and interaction mechanisms.

4. Empirical analysis

We first estimate the baseline model's parameters from US time-series data on individual and firm-level adoption, then use cross-sectional variation in demographics, firm size, and country-level adoption to validate the heterogeneous extension and identify the cross-group influence channel.

4.1. Data and sources

The data comprise six sources that jointly cover both sides of the diffusion process, individual and organisational, across the US time series and a cross-country panel.[1]

On the *individual* side, microdata from four waves of the Pew American Trends Panel (March 2023 – August 2024) provide nationally representative US adoption rates, rising from 13.6% to 33.4% [40]. Five quarterly waves of the Real-Time Population Survey (RPS), scaled to the Pew-equivalent level at the August 2024 overlap, extend the series to 41.8% by November 2025

[1] "Individual adoption" refers to personal GenAI chatbot use (Pew, RPS); "organisational adoption" refers to enterprise AI deployment (BTOS, Eurostat).

[41]. Monthly resolution between these anchor levels is provided by Google Trends search intensity for AI-related keywords, calibrated to match all four Pew adoption levels; the interpolated path supplies monthly observations for ODE estimation only and does not affect adoption levels (Appendix B.1). Eurostat's 2025 ICT Household Survey extends individual coverage to 33 European countries with demographic breakdowns [42].

On the *organisational* side, the Census BTOS tracks AI adoption bi-weekly across 19 NAICS sectors (September 2023 – October 2025), with mean sector-level adoption rising from 4.0% to 10.4% [43]. The OECD ICT Business survey [44] and Eurostat's Enterprise survey [45] extend organisational coverage to 45 countries, with country-level adoption ranging from 3.4% to 42.0% and large firms adopting at roughly three times the rate of small firms.[2]

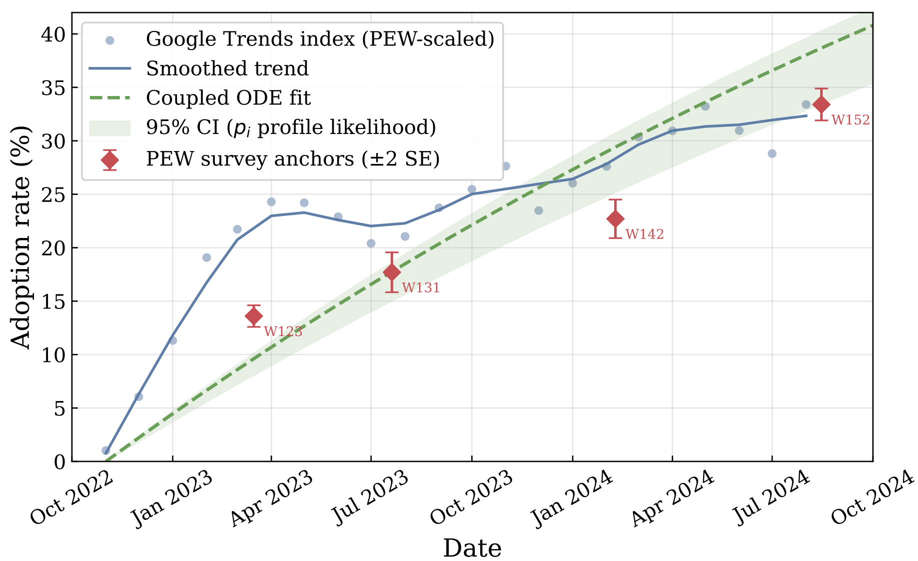

Fig. 1. Individual AI chatbot adoption in the U.S.

[2]Adoption concepts differ across sources, but parameters are estimated from each series' time profile, so cross-source level differences do not enter the estimates; see Appendix B.1 for details.

4.2. Parameter estimation

Estimation proceeds in two stages on the baseline specification (Equations 1–3). First, barrier parameters are estimated from BTOS planned-versus-actual adoption gaps, yielding $\hat{b}_0 = 0.329$ and $\hat{k} = 0.068$, with both estimates significant at $p < 0.01$. Second, with barriers fixed, we jointly estimate p_i, p_o, α, and β by minimising weighted residuals of the coupled ODE (1)–(2) against a 22-month individual adoption series[3] and 54 biweekly BTOS organisational observations; these second-stage estimates are robust to first-stage barrier uncertainty, see Appendix C.2.2 for details. Table 1 reports the estimates.

Individual and organisational adoption. The two margins diffuse through distinct mechanisms: individuals adopt through autonomous innovation, whereas the fitted dynamics attribute organisational adoption chiefly to cross-group spillover from individual uptake, gated by an eroding adoption barrier. On the individual side, the estimated innovation coefficient $\hat{p}_i = 0.023$ (Panel B) implies rapid autonomous adoption: the fitted curve rises from near zero in late 2022 to 33% by August 2024, matching three of four Pew survey anchors within two standard errors (Figure 1). On the organisational side, the joint ODE tells a mirror-image story: $\hat{p}_o$ and $\hat{\alpha}$ both settle at their lower bounds (Panel B), leaving the cross-group term $\hat{\beta}_{\text{ODE}} = 0.009$ to carry the organisational dynamics. That coefficient is not identified within-country, however—its profile likelihood is flat and the 95% interval includes zero (Appendix C.2.1)—so identification rests on the cross-country regression in Section 4.5, where multi-anchor ODE specifications recover $\hat{\beta} \in [0.010, 0.011]$ ($p < 0.004$; Table C.12). Panel C sharpens the contrast: estimating p_o sector-

[3] Google Trends search intensity calibrated to four Pew anchor levels (Appendix B.1; validated in Appendix C.1.1). Barrier and cross-group parameters are identified from BTOS sector data and cross-country variation and do not depend on the GT construction; replacing the GT shape with piecewise-linear interpolation changes $\hat{\beta}$ by less than 10%, and $\hat{p}_i$ under Pew-only estimation is 0.0201 versus 0.0228 under the baseline specification.

Table 1: Parameter estimates of the baseline model

Parameter	Estimate	SE	t-stat
Panel A: Barrier parameters (Stage 1)			
b_0	0.3290***	0.0184	17.87
k	0.0679***	0.0224	3.03
Panel B: Joint ODE (Stage 2; $q_i = q_o = 0$; $N = 76$)			
p_i	0.0228*	0.0126	1.81
p_o	0.0000[a]	0.0026	—
α	0.0000[b]	0.3052	—
β	0.0094	0.0101	0.94
Panel C: Sector-level organisational innovation			
p_o^{sector}	0.0055***	0.0011	5.21
Specification test: H_0: $q_i = q_o = 0$			
$F(2, 70) \approx 0.00$			$p = 1.00$

Notes: Panel A estimates the barrier parameters (Stage 1) by log-linear regression of the BTOS planned-versus-actual adoption gap on cumulative adoption. Panel B estimates the joint ODE (Equations 1–2, Stage 2) with b_0 and k held at their Panel A values; the restricted specification ($q_i = q_o = 0$) is adopted on the basis of the specification test (final row). Panel C reports the sector-level organisational innovation rate, estimated by NLS without cross-group terms. Standard errors are asymptotic; bootstrap confidence intervals are reported in Appendix C.2.1. [a] at the lower bound (boundary solution); [b] weakly identified, with α and p_o jointly underdetermined along a likelihood ridge. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

by-sector without the cross-group term yields $\hat{p}_o^{\text{sector}} = 0.0055$ ($p < 0.01$), an upper bound on autonomous organisational innovation because omitting βF_i loads any spillover onto p_o, and even against this upper bound $\hat{p}_i$ is significantly larger ($p = 0.013$, delta-method test), confirming that individuals innovate faster than organisations—a gap that holds despite wide sectoral dispersion, from under 5% in Construction to nearly 25% in Information, around a common upward trend.

Comparison to historical technologies. AI diffusion parameters differ markedly from prior technologies: for all historical technologies, the imitation coefficient q substantially exceeds the innovation coefficient p (typical $p/q \approx 0.03$–0.11, see Table B.10 for a survey). In the current estimation window, AI departs sharply from this pattern: the F-test confirms both imitation coefficients are indistinguishable from zero ($p = 1.00$), robust across all 19 BTOS sectors and four Pew waves. Independent behavioral evidence is consistent with this pattern: the dominant barrier reported by EU27 non-adopters in Eurostat's reasons-for-non-adoption survey is *no perceived need* (39%), while awareness and skill gaps account for only 14% combined (Appendix B.3)—suggesting that $q_i \approx 0$ reflects genuine structure rather than a window artefact alone. This contrast should nonetheless be read with caution: the truncation analysis in Section Appendix C.1.2 shows that at $T = 2$ years, Monte Carlo experiments with known imitation dynamics ($q \in [0.10, 0.50]$) fail to recover $\hat{q} > 0$ in 31–63% of replications, so $\hat{q} \approx 0$ is consistent with, but does not confirm, a genuinely innovation-driven regime.

4.3. Individual demographic heterogeneity

The extended model allows individual groups to differ in their innovation coefficients $p_{i,g}$. We test this empirically using Pew multi-wave microdata, fitting a two-parameter logistic curve[4] to each demographic group:

$$F_{g,t} = \frac{1}{1 + \left(\frac{1}{F_{g,0}} - 1\right) e^{-r_g t}}, \tag{10}$$

where r_g is the group-specific growth rate and $F_{g,0}$ the initial adoption level. Table 2 reports the estimated diffusion parameters with two emerging patterns. First, age groups differ mainly in *when* they adopt (T_{50} spans 18–43 months); growth rates are essentially uniform across working-age cohorts ($\hat{r}_g \in [0.069, 0.072]$ for 18–29, 30–49, and 50–64), with the 65+ cohort fitted at $\hat{r}_g = 0.092$ as an artefact of the early exponential phase on a low-base series. Second, education and income groups differ in both timing *and* speed: college graduates adopt faster ($\hat{r}_g = 0.077$) than those with a high-school degree or less ($\hat{r}_g = 0.057$), suggesting that educational attainment shapes not only the initial propensity to adopt but also the rate of learning—consistent with group-specific $p_{i,g}$ in the extended model.

Cross-regional validation. Comparing US and EU27 adoption reveals nearly identical aggregate levels (∼33%) and, critically, the same demographic gradients in both regions: younger and more educated individuals adopt earlier and faster (Figure 2). The consistency of these gradients across regions suggests that age and education shape adoption speed through inherent differences in digital familiarity and learning capacity, not through country-specific institutional factors—supporting the use of group-specific $p_{i,g}$ calibrated from US data for broader projections.

4.4. Firm-size heterogeneity

The extended model posits archetype-specific adoption dynamics for organisations. We estimate group-specific logistic curves (Equation 10) for

[4]Four waves cannot separately identify p and q; since the joint estimation in Table 1 finds $q \approx 0$, the logistic ($K = 1$) is the natural reduction. Section 4.6 confirms a similar Gompertz fit.

Table 2: Individual AI adoption by demographic group: Pew logistic estimates

	Group	$\hat{r}_g$	SE	T_{50}	R^2	Latest (%)
Overall		0.067	(0.006)	31.3	0.984	33.4
Age	18–29	0.070	(0.002)	17.9	0.998	54.6
	30–49	0.072	(0.010)	26.4	0.969	40.9
	50–64	0.069	(0.005)	36.4	0.991	25.3
	65+	0.092	(0.010)	42.6	0.983	11.8
Education	College+	0.077	(0.013)	21.8	0.955	49.0
	Some coll.	0.068	(0.007)	31.2	0.984	33.1
	HS or less	0.057	(0.010)	46.9	0.948	18.6
Income	Upper	0.083	(0.007)	20.8	0.989	50.0
	Middle	0.070	(0.010)	30.1	0.964	34.8
	Lower	0.057	(0.012)	41.8	0.925	23.9
Age gap		$\Delta r \approx 0.003^a$		24.7 mo	timing only	
Education gap		Δr=0.020		25.1 mo	timing & speed	
Income gap		Δr=0.026		21.0 mo	timing & speed	

Notes: $\hat{r}_g$ is the monthly logistic growth rate; T_{50} is months from the ChatGPT launch (November 2022) to 50% adoption; standard errors (in parentheses) from NLS, $n = 4$ Pew waves per group. With only four waves the estimates are descriptive—gradients rather than precise rates—but hold out-of-sample against the EU27 (Figure 2). [a]Working-age cohorts cluster at $\hat{r}_g \in [0.069, 0.072]$; the higher 65+ value reflects the early exponential phase of a low-base series, not faster adoption.

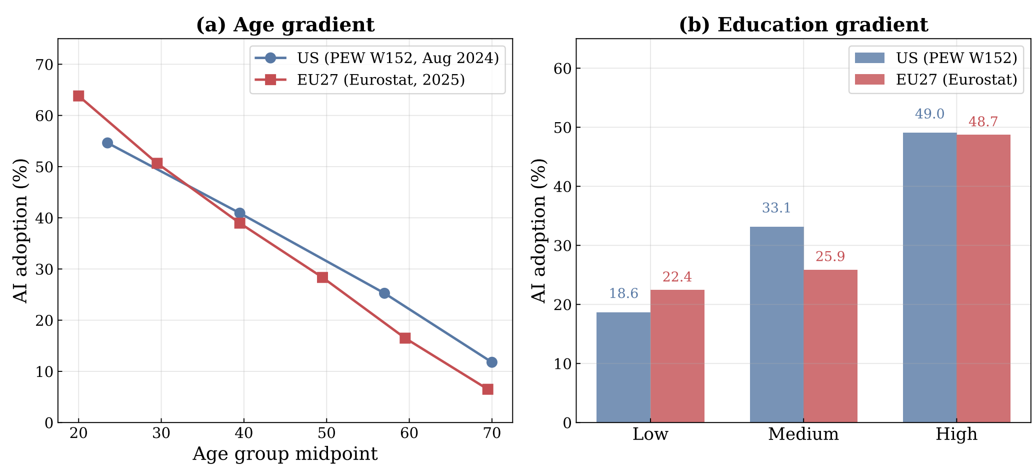

Fig. 2. Cross-regional comparison of individual AI adoption: U.S. vs. EU27.

three firm-size categories using OECD data. Table 3 reports the parameter estimates and Figure 3 shows the fitted trajectories.

Table 3: AI adoption by firm size: OECD logistic estimates

Firm size	N	$\hat{r}_h$	SE	T_{50}	R^2	Latest (%)
Small (10–49)	38	0.021	(0.007)	98	0.704	18.0
Medium (50–249)	38	0.027	(0.004)	62	0.933	30.8
Large (250+)	38	0.027	(0.004)	26	0.934	53.7
Large − Small gap	35.7 pp (mean across countries)					

Notes: Logistic curve estimates (Equation 10) fitted to the 2020–2025 OECD panel, with N the number of countries. $\hat{r}_h$ is the monthly logistic growth rate; T_{50} is months from the ChatGPT launch (November 2022) to projected 50% adoption; standard errors (in parentheses) from NLS.

The size gap is primarily one of *timing*: large firms reach 50% adoption in 26 months versus 98 for small firms, a 72-month spread, though small firms also grow modestly slower ($\hat{r}_h = 0.021$ vs. 0.027). Two reinforcing mechanisms underlie the 35.7 pp gap. Large firms face lower initial barriers, reflecting greater IT capacity and implementation resources; and they are more responsive to individual adoption spillovers, consistent with the cross-country size gradient in Section 4.5. Both patterns enter the extended model through archetype-specific $b_{0,h}$ and $\beta_{max,h}$.

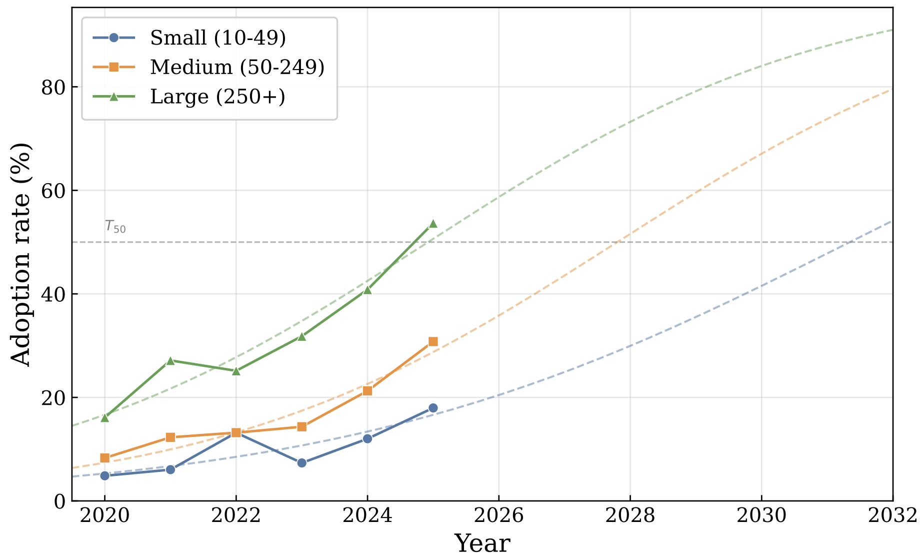

Fig. 3. Firm-size heterogeneity in AI adoption: OECD 2020–2025.

4.5. Cross-group influence estimation

An operative cross-group channel ($\beta > 0$) implies that countries with higher individual adoption should exhibit higher enterprise adoption, conditional on shared determinants. Figure 4 provides visual evidence for both model mechanisms: panel (a) shows the positive cross-country relationship between individual and enterprise AI adoption, and panel (b) shows the sector-level link between AI adoption and productivity growth.

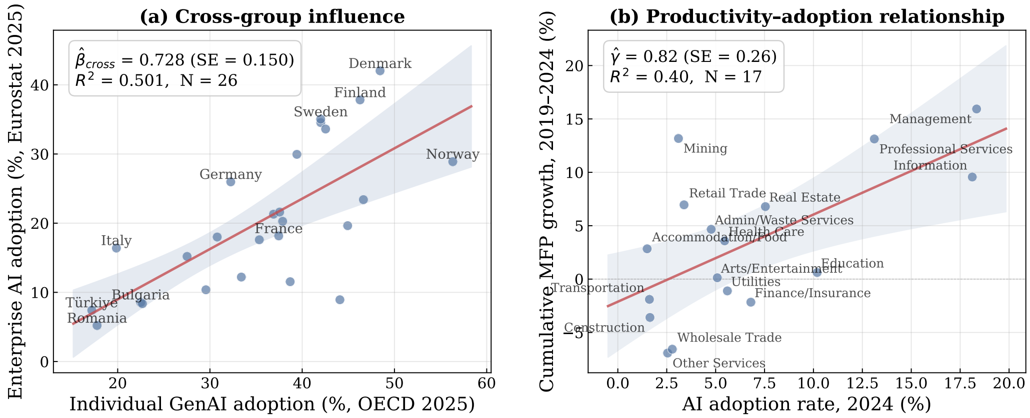

Fig. 4. Cross-sectional evidence for model mechanisms.

We test the cross-group prediction by exploiting cross-country variation in a 26-country panel, matching the 2025 OECD ICT Household Survey to enterprise AI adoption from the Eurostat ICT Enterprise Survey. We estimate:

$$F_{o,c} = \gamma_0 + \beta_{\text{cross}} F_{i,c} + \mathbf{X}_c' \delta + \varepsilon_c \tag{11}$$

where $F_{o,c}$ and $F_{i,c}$ are enterprise and individual AI adoption rates in country c, β_{cross} captures the cross-group association, and $\mathbf{X}_c$ includes controls for digital infrastructure (Going Digital score). Table 4 reports the results.

Table 4: Cross-group influence: individual GenAI and enterprise AI adoption

	(1) All firms	(2) All firms	(3) All firms	(4) Small	(5) Medium	(6) Large
Indiv. GenAI (%)	0.728***		0.378*	0.653***	1.043***	1.169***
	(0.150)		(0.178)	(0.137)	(0.193)	(0.262)
Work GenAI (%)		0.687***				
		(0.155)				
Going Digital			0.927*			
			(0.470)			
R^2	0.501	0.484	0.600	0.490	0.531	0.547
N	26	26	23	26	27	27

Notes: DV: enterprise AI adoption (%, Eurostat 2025). HC-robust SEs in parentheses. (1) bivariate; (2) work-purpose GenAI; (3) controlled for Going Digital score. (4)–(6): by firm size. Small: 10–49; Medium: 50–249; Large: 250+. $^{***}p < 0.001$, $^{*}p < 0.05$.

In the bivariate specification (Column 1), a 1 pp increase in individual GenAI adoption predicts a 0.73 pp increase in enterprise adoption ($p < 0.001$). Restricting to work-purpose GenAI (Column 2) yields a similar coefficient ($\hat{\beta}_{\text{cross}} = 0.69$). Column 3 introduces the OECD Going Digital score as a control for shared digital infrastructure—a common driver of both individual and enterprise adoption. Including this control attenuates the coefficient to $\hat{\beta}_{\text{cross}} = 0.38$ ($p = 0.033$), indicating that shared digital readiness accounts for part of the cross-country correlation; the residual captures the association that persists after netting out this common driver.[5] Columns 4–6 reveal a size gradient: $\hat{\beta}_{\text{cross}}$ rises from 0.65 (small firms) to 1.04 (medium) to 1.17 (large), showing the pattern an employee-driven spillover mechanism would predict. Panel fixed effects, instrumental variables, and a lagged specification all confirm these results, with coefficients attenuating from $\hat{\beta}_{\text{IV}} = 0.66$ to $\hat{\beta}_{\text{panel}} = 0.24$ as confounders are absorbed.[6]

Finally, to ground the productivity feedback mechanism (Eq. 9), we match BLS multifactor productivity (MFP) data at the NAICS 2-digit level [46] to BTOS sector-level AI adoption rates. Across 17 matched sectors, cumulative TFP growth over 2019–2024 is positively linked with 2024 AI adoption ($\hat{\gamma} = 0.82$, $R^2 = 0.40$, $p = 0.006$; Figure 4b), robust to outlier removal ($\hat{\gamma} = 0.70$–0.94 across leave-one-out specifications).

4.6. Robustness

We subject the core results to robustness checks across alternative functional forms, outlier treatment, sample sensitivity, and survey weighting (details in Appendix C.4).

[5] Wild bootstrap and randomisation inference confirm $p < 0.001$ for both specifications (Appendix C.2.3).

[6] With $N = 23$–27, the cross-section supports one composite control; the Going Digital score is selected as the primary observable confounder. Panel fixed effects and instrumental variables (Appendix C.2.3) absorb unobserved country-level heterogeneity.

Cross-country and survey robustness. The Going Digital–adoption regression (Table B.9) is stable under both outlier treatment—the coefficient ranges from 1.01 (trimmed tails) to 1.19 (baseline)—and leave-one-out resampling, which yields coefficients between 1.09 and 1.38 with no single country driving the result. The difference between weighted and unweighted Pew adoption rates is 3.2 pp, a minor discrepancy that does not affect qualitative conclusions.

Estimation robustness. The Gompertz alternative achieves comparable fit to the logistic ($R^2 = 0.80$ vs. 0.81); we adopt the logistic because it nests the Bass benchmark, facilitating comparison with historical parameters (Table B.10). Survey-only estimation from four Pew waves—without Google Trends—yields $\hat{p}_i = 0.020$, consistent with the monthly-series estimate of 0.023; six alternative keyword schemes produce near-identical results ($\hat{p}_i$ range: 0.001; Appendix C.1.1).

5. Calibration and simulation

Aggregate parameters (innovation rates, barriers, cross-group influence) enter the extended model directly from the empirical estimates of Section 4. Group-specific counterparts are calibrated by scaling reduced-form estimates from three cross-sectional sources: demographic adoption rates (Table 2), firm-size gradients (Table 3), and cross-country regressions (Table 4). We then simulate the model to evaluate baseline diffusion, decompose the role of heterogeneity and cross-group spillovers, and test three policy levers.

5.1. Calibration strategy

First, the barrier parameters (b_0, k) and the individual innovation coefficient $\hat{p}_i$ enter directly from Table 1. The cross-group influence parameters $\beta_{max,h}$ are calibrated from the firm-size-stratified cross-country estimates ($\hat{\beta}_{\text{cross}}^{\text{large/med/small}} = 1.17/1.04/0.65$, Table 4) by a uniform scaling factor of $\times 0.25$, yielding $\beta_{max,h} = [0.292, 0.261, 0.163]$; the factor minimises out-of-sample MAE against the held-out 2025 Eurostat wave (Appendix C.3.1). Three checks reinforce this choice: multi-anchor ODE specifications recover significant $\hat{\beta}$ ($p < 0.004$, Table C.12), alternative specifications support a positive association (Appendix C.2.3), and qualitative conclusions hold across β-scaling factors $[0.10, 0.50]$ (Appendix C.3.1). The simulation results should be interpreted as conditional on cross-group spillover operating at approximately the magnitude observed in the cross-country data. The reported simulations use the linear specialisation of the cross-group influence ($\beta_h F_i = \beta_{max,h} F_i$, $\alpha_g F_o = \alpha_{max,g} F_o$), the $\gamma \to 0$ limit of Equations 7–8, since the cross-country evidence identifies a linear individual–enterprise slope; Appendix C.3.3 recalibrates the saturating form and confirms that all qualitative conclusions are unchanged.

Second, the group-specific parameters—innovation coefficients ($p_{i,g}$, $p_{o,h}$) and cross-group influence parameters ($\beta_{max,h}$, $\alpha_{max,g}$)—are obtained by scaling reduced-form estimates from the cohort, firm-size, and cross-country regressions (Tables 2, 3, and 4).[7] Scaling factors are chosen to minimise in-sample MAE against Pew and Eurostat adoption, and all simulation conclusions are stress-tested for sensitivity to this choice.

Third, the productivity feedback parameters (φ, λ, P_{max}, δ) are set from the literature [13, 39], with φ anchored to the BLS productivity–adoption slope (Figure 4b) and subjected to sensitivity analysis in Appendix D.4. The complete parameterisation, with baseline values and calibration sources for each parameter block, is summarised in Table D.24.

5.2. Simulation results and calibration fit

Baseline fit and out-of-sample checks. Figure 5 overlays the calibrated model on observed data: the model reproduces the Eurostat enterprise series within

[7] Derivations, scaling values, and sensitivity are reported in Appendix C.3.1 and Appendix C.3.1.

0.7 pp at every observed year and tracks the four Pew individual waves with a mean absolute error of 3.4 pp (per-wave residuals 2–6 pp), for a combined in-sample MAE of 2.17 pp (Table C.18).[8] The simulated individual trajectory runs above the observed Pew levels because the coupled model layers cross-group spillover and productivity feedback onto the single-series logistic fit, accelerating simulated uptake. Under the benchmark calibration, the model projects 93.6% individual and 74.1% organisational adoption by 2033, with the productivity index reaching $P(10) = 1.49$.[9]

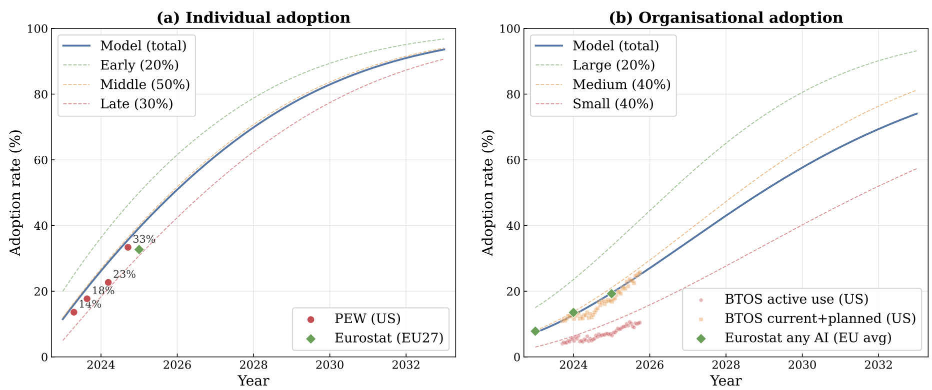

Fig. 5. Simulated individual and organisational adoption against observed data, 2023–2033.

Key drivers. A ±20% parameter sensitivity analysis (Figure D.7) reveals that adoption speed is driven primarily by barrier height ($b_{0,h}$), erosion rate (k_h), and cross-group influence ($\beta_{max,h}$), while terminal productivity depends on the adoption-to-productivity efficiency φ and organisational feedback sensitivity $\lambda_{o,h}$. This distinction—barriers and spillovers drive *adoption pace*,

[8]The in-sample fit is computed against US Pew waves; the Eurostat comparison is an out-of-sample cross-check against independently collected European data.

[9]$P(10) = 1.49$ is the baseline ODE at year 10; the Monte Carlo mean across 1,000 parameter draws is $\bar{P}_{final} = 1.43$ with 95% interval $[1.05, 1.73]$ (Table D.26).

while φ and λ drive *economic impact*—motivates the policy experiments below.

Decomposing model mechanisms. To quantify the contribution of each model extension, we run three counterfactual simulations that progressively disable features (full results in Appendix D.2). First, disabling the cross-group channels ($\alpha = \beta = 0$) is the single most consequential change: organisational adoption never reaches 50% within the ten-year horizon (it saturates near 36%, against the calibrated baseline's 74%), and terminal productivity falls by 14%. In the calibrated model, cross-group spillover from individual to enterprise adoption is therefore the dominant accelerator of organisational uptake within the modelled horizon. Second, collapsing to homogeneous groups ($G = M = 1$) moves *aggregate* organisational T_{50} by -0.16 to $+0.64$ years across two homogeneous comparators—a small premium, consistent with the approximate-aggregation property familiar from heterogeneous-agent macro [47]. The within-population timing dispersion is much larger: the heterogeneous model generates a 5.2-year T_{50} spread across firm-size archetypes (Large 3.5 yr, Medium 5.3 yr, Small 8.7 yr), capturing most of the 6-year empirical spread in Table 3. Heterogeneity does not materially accelerate aggregate timing but resolves a within-population gradient that aggregate models cannot represent. Third, disabling the productivity feedback loop delays organisational T_{50} by approximately 0.2 years and reduces terminal productivity by 0.5% at baseline λ, but by 1.0 years and 4.2% when λ is amplified to isolate the mechanism, consistent with self-reinforcing growth dynamics [48]. Cross-group spillover dominates the adoption-timing dimension; heterogeneity dominates the composition dimension; the productivity-feedback loop links adoption to economic outcomes with a magnitude that depends on the assumed feedback strength.

5.3. Policy experiments

Motivated by the policy evidence in Stylized Fact 6, we evaluate three direct policy levers against a calibrated baseline simulation with parameter values reported in Table D.24. *Direct Subsidies* reduce organisational barriers, while *Education and Training* works on two channels at once, lifting individual innovation and easing organisational barriers in tandem. *R&D Incentives*, by contrast, target the productivity-feedback channel, strengthening the loop linking adoption to productivity gains.

To compare scenarios on a common scale, we define cost as intervention intensity — the sum of absolute percentage changes in the targeted parameters [49] — and effectiveness as adoption gain in percentage points per unit of cost. Because cost reflects parameter adjustment rather than fiscal expenditure, these are comparative model exercises rather than welfare evaluations. Table 5 reports adoption gains, cost metrics, and effectiveness ratios by scenario.

Table 5: Policy effectiveness with the empirically calibrated model

	5-Year Horizon (2028)			10-Year Horizon (2033)		
Policy	ΔF_i	ΔF_o	BCR	ΔF_i	ΔF_o	BCR
Direct Subsidies	+0.8	+4.4	0.17	+0.5	+5.2	0.19
Education & Training	+5.0	+3.8	0.25	+2.0	+4.0	0.17
R&D Incentives	+0.0	+0.3	0.01	+0.1	+1.1	0.02

Notes: 5- and 10-year horizons measured from the 2023 simulation start, following ChatGPT's late-2022 release. ΔF_i and ΔF_o are percentage-point changes relative to baseline. BCR = benefit-cost ratio (total adoption gain in pp per cost unit). Model calibrated to Pew (individual) and OECD/Eurostat (enterprise) observed adoption data.

The results reveal a clear ranking. *Education and Training* is the most effective intervention per unit of intensity (effectiveness ratio = 0.25), achieving a 5.0 pp individual and 3.8 pp organisational adoption gain by 2028. *Direct Subsidies* (ratio = 0.17) accelerate organisational adoption (+4.4 pp) but have minimal impact on individual adoption. *R&D Incentives*, despite the highest intervention intensity, produce the smallest adoption gains (ratio = 0.01), suggesting that amplifying the productivity feedback loop has limited effect during the early diffusion phase when the adoption base is still small. The policies also differ in their distributional effects. Under all scenarios, the Large−Small firm adoption gap rises through the mid-2020s as large firms adopt first, peaks around 2030, and then narrows as smaller firms catch up. *Direct Subsidies* is the most size-equalising intervention, narrowing the 2033 gap by 7.8 pp relative to baseline because small firms—which start with higher barriers—benefit proportionally more from a uniform barrier cut; *Education and Training* narrows it by 4.9 pp, while *R&D Incentives* leaves the gap essentially unchanged.

6. Conclusion

This paper develops an extended Bass diffusion model for heterogeneous AI adoption across individuals and organisations. The model incorporates four features absent from the standard framework: heterogeneous adopter archetypes, bidirectional cross-group influences, endogenously eroding adoption barriers, and productivity feedback. Its base dynamics are estimated from survey data on individual and organisational adoption; the full extended model is calibrated from auxiliary cross-sectional evidence and then simulated over a ten-year horizon.

Two empirical results stand out. First, cross-country evidence is consistent with cross-group spillover — rather than within-group imitation — as the dominant accelerator of organisational AI adoption. The estimated imitation coefficients are indistinguishable from zero in the estimation window, though a two-year span may be too short to detect imitation even where it operates; social influence instead operates between populations, with individual adoption strongly associated with enterprise adoption across countries. The cross-country coefficient of 0.73 is robust to panel fixed-effects, instrumental-variable, and lagged specifications, although full causal identification awaits firm-level microdata or natural-experiment designs. The channel is quantitatively decisive: switching off cross-group influence in the calibrated model leaves organisational adoption saturating near 36%, short of the 50% threshold even after ten years. Second, heterogeneity earns its complexity not by accelerating the aggregate trajectory but by resolving within-population structure that homogeneous models smooth over. Aggregate organisational timing shifts by only -0.16 to $+0.64$ years when the heterogeneous populations are replaced with homogeneous comparators, yet within the heterogeneous model the time to 50% adoption ranges from 3.5 years for large firms to 8.7 years for small ones—a 5.2-year spread that captures most of the six-year gap in the OECD estimates and underlies the 35.7 pp firm-size adoption gap, the demographic gradient in individual uptake, and the differential effectiveness of group-targeted policies. The same gradients appear in US and European data alike, indicating that the underlying mechanisms are structural rather than country-specific.

The calibrated model tracks the observed individual and enterprise series with a combined in-sample mean absolute error of 2.17 percentage points, and simulates overall organisational adoption reaching 50% by 2029 and 74% by 2033 under the benchmark calibration, with diffusion translating into sustained productivity growth over the horizon. In the policy experiments, Education and Training delivers the largest adoption gain per unit of intervention intensity, while Direct Subsidies are the most size-equalising, narrowing the 2033 gap between large and small firms by 7.8 pp. These findings recommend treating AI diffusion as an ecosystem of interacting heterogeneous agents, where effective policy targets the connections between individual and organisational adoption.

Several limitations qualify these results. Structurally, adoption is irreversible, and the model is demand-driven, abstracting from supply-side constraints and from strategic interaction among firms. Some organisational barriers may be permanent rather than transitory, though a permanent barrier floor consistent with the BTOS evidence reduces ten-year organisational adoption by only 0.6 percentage points. Empirically, the short 2023–2025 window limits identification of the imitation coefficient and the time-series cross-group parameter, and the productivity-feedback parameters are calibrated from the literature rather than estimated. The extended model is itself calibrated rather than fully estimated, so its projections are conditional on the mapping assumptions of Section 5.1, with organisational T_{50} spanning 4.8 to 8.3 years and projected 2030 adoption spanning 41 to 75 percentage points across plausible calibration choices. Future work should pursue matched employer–employee data for causal cross-group inference, longer panels to resolve imitation dynamics, and sector-level estimation for more targeted policy analysis.

Funding

This work was supported by the French government through the France 2030 investment plan managed by the National Research Agency (ANR), as part of the Initiative of Excellence Université Côte d'Azur under reference number ANR-15-IDEX-01.

Declaration of generative AI and AI-assisted technologies in the manuscript preparation process

During the preparation of this work, the authors used Claude (Anthropic) to assist with developing and debugging the simulation and estimation code and to improve the readability and language of the manuscript, alongside basic grammar and spell-checking tools. The research design, model specification, empirical strategy, and interpretation of results are the authors' own. After using these tools, the authors reviewed, tested, and validated all AI-assisted code, reviewed and edited the text as needed, and take full responsibility for the content of the published article.

Data Availability

The simulation code will be made available in a public repository upon publication. The empirical analysis uses the following datasets: Pew Research Center American Trends Panel Waves 123, 131, 142, and 152 (March 2023 – August 2024), available at `https://www.pewresearch.org/` (free registration required; microdata distributed under Pew's terms of use); Real-Time Population Survey generative AI adoption series, published aggregates available at `https://genaiadoptiontracker.com/` (microdata via OpenICPSR project 158081); US Census Bureau Business Trends and Outlook Survey (BTOS), available at `https://www.census.gov/`; OECD ICT Access and Usage by Businesses, available via the OECD SDMX API; Eurostat Community Survey on ICT Usage in Enterprises (`isoc_eb_ai`, reference year 2024, published December 2024) and ICT Household Survey on Individual AI Use (`isoc_ai_iaiu`, reference year 2025, published January 2026), available at `https://ec.europa.eu/eurostat/`. Processed datasets and estimation scripts will be included in the repository.

References

[1] C. Hu, S. Chan, Chatgpt sets record for fastest-growing user base - analyst note, *Reuters* (2023). Accessed on June 11, 2025.

[2] McKinsey & Company, The state of AI in 2025: Agents, innovation, and transformation, 2025. Accessed November 6, 2025.

[3] F. M. Bass, A new product growth for model consumer durables, *Management science* 15 (1969) 215–227.

[4] Z. Griliches, Hybrid corn: An exploration in the economics of technological change, *Econometrica* 25 (1957) 501–522.

[5] S. Davies, The Diffusion of Process Innovations, Cambridge University Press, Cambridge, 1979.

[6] M. Karshenas, P. L. Stoneman, Rank, stock, order, and epidemic effects in the diffusion of new process technologies: An empirical model, *the RAND Journal of Economics* (1993) 503–528.

[7] G. Battisti, P. Stoneman, The intra-firm diffusion of new technologies: Profiling the non-adopter, *International Journal of Industrial Organization* 21 (2003) 1419–1444.

[8] G. Battisti, P. Stoneman, The intra-firm diffusion of complementary innovations: An empirical analysis, *Research Policy* 34 (2005) 1–9.

[9] J. A. Norton, F. M. Bass, A diffusion theory model of adoption and substitution for successive generations of high-technology products, *Management science* 33 (1987) 1069–1086.

[10] G. D. Prato, M. L. Cobo, S. Samoili, R. Righi, M. V. P. Baillet, M. Cardona, The AI Techno-Economic Segment Analysis, JRC Research Reports JRC118071, Joint Research Centre, 2019.

[11] J. S. Gans, AI Adoption in a Competitive Market, Working Paper 29996, National Bureau of Economic Research, 2022.

[12] X. Ma, J. Wang, Simulation analysis of artificial intelligence technology diffusion under market competition and policy incentives based on complex network evolutionary game models, *Systems* 12 (2024) 242.

[13] T. Gries, W. Naudé, Artificial intelligence, jobs, inequality and productivity: Does aggregate demand matter?, *SSRN Electronic Journal* (2018).

[14] T. Gries, W. Naudé, Artificial Intelligence, Income Distribution and Economic Growth, Technical Report, Institute of Labor Economics (IZA), 2020.

[15] B. Ünveren, T. Durmaz, S. Sunal, Ai revolution and coordination failure: Theory and evidence, *Journal of Macroeconomics* 78 (2023) 103561.

[16] M. Comunale, The economic impacts and the regulation of ai: A review of the academic literature and policy actions, *IMF Working Papers* 2024 (2024) 1.

[17] G. Silverberg, G. Dosi, L. Orsenigo, Innovation and the dynamics of the firm: A new approach, *Journal of Economic Behavior & Organization* 10 (1988) 35–58.

[18] H. Dawid, Agent-based models of innovation and technological change, in: L. Tesfatsion, K. L. Judd (Eds.), Handbook of Computational Economics, volume 2, Elsevier, 2006, pp. 1235–1272.

[19] J.-C. Rochet, J. Tirole, Platform competition in two-sided markets, *Journal*

of the European Economic Association 1 (2003) 990–1029.

[20] T. F. Bresnahan, M. Trajtenberg, General purpose technologies: 'engines of growth'?, *Journal of Econometrics* 65 (1995) 83–108.

[21] E. Helpman, M. Trajtenberg, Diffusion of general purpose technologies, *NBER Working Paper Series* (1996).

[22] P. A. David, G. Wright, General purpose technologies and surges in productivity: Historical reflections on the future of the ict revolution, in: S. Dutta, B. Bilbao-Osorio (Eds.), The Global Information Technology Report 2011, World Economic Forum, 2012, pp. 23–42.

[23] R. A. Lipcsey, Ai diffusion to low- and middle income countries; a blessing or a curse?, 2025. `arXiv:2405.20399`.

[24] I. S. Gill, W. Fengler, K. Karakulah, The economics of ai-based technologies: A framework and an application to europe, *SSRN Electronic Journal* (2020).

[25] J. Vithlani, C. Hawksworth, J. Elvidge, L. Ayiku, D. Dawoud, Economic evaluations of artificial intelligence-based healthcare interventions: a systematic literature review of best practices in their conduct and reporting, *Frontiers in Pharmacology* 14 (2023).

[26] S. Somjai, K. Jermsittiparsert, T. Chankoson, Determining the initial and subsequent impact of artificial intelligence adoption on economy: a macroeconomic survey from asean, *Journal of Intelligent & Fuzzy Systems* 39 (2020) 5459–5474.

[27] A. Upreti, V. Sridhar, Artificial intelligence and its effect on employment and skilling, Routledge India, 2021, pp. 31–55.

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

[29] S. Min, B. Kim, Ai technology adoption in corporate it network operations based on the toe model, *Digital* 4 (2024) 947–970.

[30] J. B. O'Donnell, Auditing transformation: A model of artificial intelligence adoption, *Journal of Applied Business and Economics* 26 (2024).

[31] L. Abrardi, C. Cambini, L. Rondi, Artificial intelligence, firms and consumer behavior: A survey, *Journal of Economic Surveys* 36 (2021) 969–991.

[32] A. Korinek, J. Stiglitz, Artificial Intelligence and Its Implications for Income Distribution and Unemployment, Technical Report, National Bureau of Economic Research, 2017.

[33] B. Alnasser, A review of literature on the economic implications of implementing artificial intelligence in healthcare, *E-Health Telecommunication Systems and Networks* 12 (2023) 35–48.

[34] B. P. E. Research, Productivity growth and employment: Ai literature review, *BNP Paribas Economic Research* (2025). Accessed: 2025-01-06.

[35] OECD, Oecd employment outlook 2023: Artificial intelligence, job quality and inclusiveness, *OECD Employment Outlook* (2023). Accessed: 2025-01-06.

[36] S. I. of Artificial Intelligence, Ai productivity in education: Real gains, costs, and what to do next, *SIAI Research* (2025). Accessed: 2025-01-06.

[37] R. Sun, D. Trefler, The impact of ai and cross-border data regulation on international trade in digital services: A large language model, *SSRN Electronic Journal* (2023).

[38] J. Wolff, J. Pauling, A. Keck, J. Baumbach, Systematic review of economic impact studies of artificial intelligence in health care, *Journal of Medical Internet Research* 22 (2020) e16866.

[39] B. H. Hall, J. Mairesse, P. Mohnen, Measuring the returns to R&D, *Handbook of the Economics of Innovation* 2 (2010) 1033–1082.

[40] Pew Research Center, [dataset] American Trends Panel, Waves 123, 131, 142, and 152, https://www.pewresearch.org/, 2024.

[41] A. Bick, A. Blandin, D. J. Deming, [dataset] Real-Time Population Survey, Generative AI Adoption Tracker, https://genaiadoptiontracker.com/, 2025. Microdata available via OpenICPSR project 158081.

[42] Eurostat, [dataset] Individuals – Use of Generative Artificial Intelligence Tools (isoc_ai_iaiu), https://ec.europa.eu/eurostat/, 2025.

[43] US Census Bureau, [dataset] Business Trends and Outlook Survey (BTOS), https://www.census.gov/, 2025.

[44] OECD, [dataset] ICT Access and Usage by Businesses (SDMX dataflow DSD_ICT_B@DF_BUSINESSES), OECD SDMX API, 2025.

[45] Eurostat, [dataset] Community Survey on ICT Usage and E-commerce in Enterprises (isoc_eb_ai, indicator E_AI_TANY), https://ec.europa.eu/eurostat/, 2024.

[46] Bureau of Labor Statistics, [dataset] Total Factor Productivity for Major Industries (KLEMS), https://www.bls.gov/productivity/, 2025.

[47] P. Krusell, A. A. Smith, Jr., Income and wealth heterogeneity in the macroeconomy, *Journal of Political Economy* 106 (1998) 867–896.

[48] P. M. Romer, Endogenous technological change, *Journal of Political Economy* 98 (1990) S71–S102.

[49] P. Stoneman, The Economics of Technological Diffusion, Blackwell Publishers, Oxford, 2002.

[50] R. Chatterjee, J. Eliashberg, The innovation diffusion process in a heterogeneous population: A micromodeling approach, *Management Science* 36 (1990) 1057–1079.

[51] H. P. Young, Innovation diffusion in heterogeneous populations: Contagion, social influence, and social learning, *American Economic Review* 99 (2009) 1899–1924.

[52] E. M. Rogers, Diffusion of innovations, 5th ed., Free Press, 2003.

[53] V. Mahajan, E. Muller, F. M. Bass, New product diffusion models in marketing: A review and directions for research, *Journal of Marketing* 54 (1990) 1–26.

[54] OECD, [dataset] Going Digital Toolkit, https://goingdigital.oecd.org/, 2025.

[55] A. Bick, A. Blandin, D. J. Deming, The rapid adoption of generative AI, *Management Science* (2025). NBER Working Paper No. 32966.

[56] N. Meade, T. Islam, Modelling and forecasting the diffusion of innovation – a 25-year review, *International Journal of Forecasting* 22 (2006) 519–545.

[57] YouGov, What Americans Think About ChatGPT and AI Text, Technical Report, YouGov, 2023. Survey of 1,000 US adults, January 24–27, 2023. Non-probability sample matching.

[58] Ipsos, Americans Hold Mixed Opinions on AI and Fear Its Potential to Disrupt Society, Technical Report, Ipsos, 2023. KnowledgePanel (probability-based,

address-based sampling), N = 1,008, April 20–24, 2023.

[59] Consumer Reports, NORC at the University of Chicago, AI Chatbot Surveys: How Americans Are Using and Thinking About Text-Based Generative AI, Technical Report, Consumer Reports, 2024. AmeriSpeak probability panel, N = 2,062, field through August 21, 2023. Published January 31, 2024.

[60] C. F. Manski, Identification of endogenous social effects: The reflection problem, *The Review of Economic Studies* 60 (1993) 531–542.

[61] R. Davidson, E. Flachaire, The wild bootstrap, tamed at last, *Journal of Econometrics* 146 (2008) 162–169.

[62] J. G. MacKinnon, M. D. Webb, Wild bootstrap inference for wildly different cluster sizes, *Journal of Applied Econometrics* 32 (2017) 233–254.

[63] G. Kaplan, B. Moll, G. L. Violante, Monetary policy according to HANK, *American Economic Review* 108 (2018) 697–743.

Appendix A. Modeling details

Appendix A.1. Microfoundations of the diffusion equations

This appendix derives the reduced-form diffusion equations of Section 3.1 from an explicit model of heterogeneous adopters, showing that the Bass coefficients (p, q, α, β) and the adoption barrier aggregate distinct micro-level channels rather than being free parameters. The mapping justifies the linear hazard as a consequence of distributional assumptions and identifies the economic mechanisms each estimated coefficient summarises.

Threshold model. The adoption equations follow from a heterogeneous-cost threshold model.[10] Each not-yet-adopted agent j of type $\tau \in \{i, o\}$ faces, per unit time, an adoption cost c_j drawn from a distribution G_τ on $[0, \bar{c}_\tau]$ and

[10] Bayesian social learning [50] and evolutionary dynamics [51] yield the same adoption-proportional hazard; we use the threshold model because it accommodates cross-group channels.

adopts when net benefit $V_\tau(t) \geq c_j$:

$$V_i(t) \;=\; \underbrace{v_0}_{\text{autonomous}} \;+\; \underbrace{v_1\, F_{i,t}}_{\substack{\text{within-group}\\\text{learning}}} \;+\; \underbrace{v_2\, F_{o,t}}_{\substack{\text{cross-group}\\\text{signal}}}\;, \tag{A.1}$$

with an analogous form for organisations, $V_o(t) = w_0 + w_1 F_{o,t} + w_2 F_{i,t} + w_3(P_t - 1)$, where $w_3(P_t - 1)$ adds the productivity-gain signal of the extended model (zero at the pre-AI baseline $P_0 = 1$). Under uniformly distributed costs, $G_\tau(c) = c/\bar{c}_\tau$, a non-adopter's benefit clears its cost draw with probability $G_\tau(V_\tau(t)) = V_\tau(t)/\bar{c}_\tau$—its adoption hazard—so the flow of new adopters is this hazard applied to the remaining non-adopter pool $(1 - F_{\tau,t})$:

$$\frac{dF_{\tau,t}}{dt} = [p_\tau + q_\tau F_{\tau,t} + \alpha_\tau F_{-\tau,t}]\,(1 - F_{\tau,t}), \tag{A.2}$$

where $F_{-\tau,t}$ denotes the adoption level of the other population, and $p_\tau \equiv v_0/\bar{c}_\tau$, $q_\tau \equiv v_1/\bar{c}_\tau$, $\alpha_\tau \equiv v_2/\bar{c}_\tau$—recovering the Bass specification augmented with cross-group influence (Equations 1–2). The cross-group coefficient $\beta \equiv w_2/\bar{c}_o$ captures the marginal effect of individual adoption on organisational net benefit, scaled by cost heterogeneity. Because organisations face greater cost dispersion ($\bar{c}_o \gg \bar{c}_i$), $p_o < p_i$, consistent with slower organisational adoption. The uniform distribution is a tractability assumption: log-normal costs yield probit diffusion [5], exponential costs yield Gompertz dynamics. The linear hazard is therefore a consequence of distributional choice, not an axiom; we adopt it because it preserves analytical tractability while nesting the standard Bass model as a special case ($v_2 = 0$).

Adoption barrier. For organisations, the adoption barrier b_t enters as a proportional cost shifter: the effective threshold becomes $c_j/(1 - b_t)$, raising the hurdle for all firms while preserving cross-firm cost heterogeneity. This yields the organisational hazard in Equation 2. The barrier declines as adoption accumulates, following the exponential form $b_t = b_0 \exp(-k(F_{i,t} + F_{o,t}))$ in

Equation 3. We prefer this to a logistic barrier $b_0/(1 + e^{k(A-c)})$, which would hold barriers high until a threshold adoption level is reached; the exponential form instead lets barriers fall from the very first adopters—appropriate for AI, where early uptake already clarifies use cases, surfaces best practices, and seeds vendor ecosystems [6].

Economic channels underlying α and β. The forward channel β aggregates four mechanisms. Individual AI use builds human capital that lowers organisational training costs [28]; employee productivity gains reduce uncertainty about returns [52]; firms formalise AI use once their workforce already employs it informally [8]; and widespread interface familiarity lowers integration costs [20]. The reverse channel α operates symmetrically, through workplace exposure at AI-adopting firms [7], legitimisation by formal deployment [53], and infrastructure spillovers from organisational support ecosystems [22].

All channels are consistent with the linear specification in the base model: each increases $V_\tau(t)$ proportionally to the other group's adoption level. However, the channels differ in timing and saturation behaviour—workforce readiness accumulates gradually while competitive pressure may exhibit threshold effects—motivating the nonlinear generalisation in Section 3.2, exercised as a robustness check in Appendix C.3.3. The reduced-form specification is agnostic about the relative contribution of each channel; our empirical estimate of β captures their aggregate effect. This distinction is relevant for policy design: education interventions target the workforce-readiness channel, while information-disclosure policies target the demonstrated-viability channel (Section 5.3).

Appendix A.2. Model validation

Three internal consistency checks verify that the extended model reduces correctly to its base case, that its productivity and adoption blocks are modular, and that it generates no spontaneous adoption in the absence of diffusion drivers.

Homogeneous restriction equivalence. Under the homogeneous restriction ($G = M = 1$) with productivity feedback disabled, the extended model reproduces the base two-population specification to within numerical tolerance: fitting identical parameters (p_i, q_i, p_o, q_o, α, β) yields $R^2 > 0.999$ for both $F_{i,t}$ and $F_{o,t}$ trajectories relative to the base-case solution.

Productivity ODE neutrality. Setting all productivity parameters to zero ($\varphi = 0$, $\delta = 0$) while leaving the full heterogeneous adoption structure active holds the productivity level constant at its initial value and leaves the adoption trajectories unchanged relative to a baseline run without the productivity block—maximum deviation below 10^{-9}. The productivity and adoption modules therefore do not contaminate each other numerically.

Boundary condition analysis. Setting all innovation, imitation, and cross-group parameters to zero ($p_{i,g} = q_{i,g} = \alpha_{\max,g} = p_{o,h} = q_{o,h} = \beta_{\max,h} = 0$) leaves $F_{i,t}$ and $F_{o,t}$ pinned at their initial near-zero values throughout the simulation horizon, confirming that no spontaneous adoption is generated in the absence of diffusion drivers.

Appendix A.3. Formal steady-state and stability analysis

A closed-form analytical steady state is intractable due to the coupled non-linear dynamics, so we verify the equilibrium numerically and assess local stability via the Jacobian.

Numerical steady-state verification. Under baseline parameterisation, a long-horizon simulation ($T = 500$) converges to full adoption with elevated productivity: $F_i^{ss} = F_o^{ss} = 1.000$ and $P^{ss} = 1.881$. The productivity level admits a closed form once adoption saturates: setting $dP/dt = 0$ with $F_i = F_o = 1$ gives $P^{ss} = \varphi(s_i + s_o)P_{\max}/[\delta + \varphi(s_i + s_o)]$, which evaluates to 1.881 at baseline. The system therefore admits a unique attractor at the corner of the unit simplex, consistent with the theoretical result that any positive innovation coefficient guarantees eventual full adoption in the absence of a permanent barrier (Appendix C.3.2). A potential concern is that the productivity feedback loop could sustain a stable low-adoption equilibrium—groups that start behind gain less productivity, which further slows their adoption—but the calibrated innovation coefficients $p_{i,g}$ and $p_{o,h}$ are large enough to pull all groups toward full adoption regardless of initial conditions; productivity feedback therefore only accelerates convergence rather than gating it. Global convergence to this equilibrium follows from monotonicity: any positive innovation coefficient gives $dF/dt > 0$ below unity, so all interior initial conditions are drawn to it.

Stability analysis. The Jacobian matrix J evaluated at the steady state has all eigenvalues with strictly negative real parts across $\varphi \in [0.1, 2.0]$ (other parameters held at baseline), confirming that the full-adoption equilibrium is locally stable throughout the relevant parameter space.

Appendix B. Data and estimation

This appendix reports data sources, descriptive statistics, and the full set of sector-level and country-level parameter estimates underlying the analysis in Section 4.

Appendix B.1. Data sources and construction

Table B.6 reports descriptive statistics for the five datasets used in the empirical analysis; the paragraphs that follow document construction choices for each source.

Pew American Trends Panel. We use four Pew ATP waves: W123 (March 2023, $N = 10{,}701$; variables `GPT2_a/b/c`), W131 (July 2023; `GPTUSE`), W142 (February 2024; `GPTUSE`), and W152 (August 2024, $N = 5{,}410$; broadened to "AI chatbots such as ChatGPT, Gemini, or Copilot"). W131 and W142 use a form split: only Form 1 receives the adoption question, yielding effective samples of $N = 2{,}532$ and $N = 5{,}050$. All rates apply the wave-specific design weights and recover the published rates within 0.4 pp (13.6%, 17.7%, 22.7%, 33.4%).

Table B.6: Data sources and descriptive statistics

Variable	N	Mean	SD	Min	Max
Panel A: Pew American Trends Panel [40]; US individuals; W152, Aug. 2024					
Used AI chatbot (binary)	5,410	0.37	0.48	0.00	1.00
Used any AI tool (binary)	5,377	0.60	0.49	0.00	1.00
Barrier perception index	5,401	2.13	0.77	1.00	5.00
Age 18–29	5,410	0.14	0.35	0.00	1.00
Bachelor's degree or higher	5,410	0.43	0.49	0.00	1.00
Upper income tier	5,079	0.22	0.42	0.00	1.00
Panel B: US Census BTOS [43]; US firms (19 NAICS sectors); Sep. 2023 – Oct. 2025					
Current AI adoption (%)	881	6.93	5.65	0.70	29.80
Planned AI adoption (%)	895	9.65	7.08	1.40	39.80
Adoption by firm size (%)	162	5.79	1.96	2.50	11.00
Panel C: OECD ICT Business [44]; 38 countries; 2020–2025					
All firms, 10+ employees (%)	154	11.17	8.46	1.38	42.03
Small firms, 10–49 (%)	154	9.39	7.56	1.06	37.45
Large firms, 250+ (%)	154	32.99	17.74	7.13	79.40
Panel D: Cross-country covariates (45 countries, latest year)					
Latest AI adoption (%)	45	16.44	10.58	3.37	42.03
OECD Going Digital score [54]	35	64.52	5.22	54.80	76.10
R&D expenditure (% of GDP)	28	1.72	0.89	0.46	3.40
Panel E: Eurostat Individual and Enterprise [42, 45]; 33 countries; 2025					
Individual GenAI adoption (%)	33	35.62	10.56	17.19	56.32
Work AI use (%)	33	17.20	8.36	1.31	35.44
Matched enterprise AI adoption (%)	31	19.60	9.79	5.21	42.03

Notes: N is the number of respondents (Panel A), sector×period observations (Panel B), country×year observations (Panel C), or countries (Panels D–E). Panel A pools four Pew waves (W123, W131, W142, W152; ∼5,400 respondents each under probability-based sampling), with the statistics shown drawn from the most recent wave (W152, August 2024). Panel B (BTOS) is restricted to periods 31–84, which share consistent question wording at biweekly frequency.

Real-Time Population Survey (RPS). Five quarterly RPS waves [41] extend the US individual series from August 2024 through November 2025 (∼5,000–6,000 respondents per wave; US adults aged 18–64). RPS asks about generative AI use broadly rather than chatbot use specifically [55], so its rates are systematically higher (44.6% vs. Pew W152's 33.4% at the August 2024 overlap). We rescale each RPS rate by the level-shift factor $33.4/44.6 = 0.748$ to match the Pew scale at the overlap, yielding a nine-point combined series (4 Pew + 5 scaled RPS, March 2023 – November 2025) used in robustness checks (Appendix C.2.1, Table C.14, Panel B). Baseline estimation uses Pew-only to preserve within-source scope consistency.

Census BTOS. BTOS collects data bi-weekly across 19 NAICS sectors. We restrict the sample to periods 31–84 (September 2023 – October 2025), which use a consistent question asking whether the firm used AI "in producing goods or services." Periods 88 onward broaden to "any business functions," creating a structural break; the restriction preserves temporal comparability.

Reconciling adoption concepts across sources. The adoption series use three distinct concepts: Pew (individual chatbot use), BTOS (firms' current AI use in production), and Eurostat Enterprise (firms' use of any AI technology). We address scope differences in two ways. First, question wording is held constant within each series, so the estimated parameters (p, q, b_0, k) reflect within-source growth rather than scope shifts. Second, BTOS's current-plus-planned aggregate aligns closely with Eurostat's broader "any AI" rate across countries (Figure 5, panel b)—the two capture the same diffusion process at different points on the use-curve (intention → implementation). The model is calibrated on the internally consistent US measures (Pew chatbot, BTOS current-use) and validated out-of-sample against Eurostat EU27.

Individual adoption time series. Four Pew snapshots pin down adoption *levels* but cannot identify trajectory *shape* and *speed*, which the coupled-ODE estimation requires alongside the biweekly organisational series. We therefore construct a monthly series from Google Trends search intensity, using a weighted composite of seven keywords (ChatGPT 50%, Claude AI 15%, generative AI 15%, AI chatbot 10%, LLM 5%, GitHub Copilot 2.5%, Midjourney 2.5%) whose weights approximate US search-volume shares during the estimation window. The composite is rescaled so its August 2024 value matches the Pew W152 rate of 33.4%, producing a 22-month series (November 2022 – August 2024) that supplies the temporal granularity needed to estimate p_i, q_i and to identify the cross-group channels α, β (Equations 1–2). Validation against three withheld Pew waves and against a survey-only re-estimation ($\hat{p}_i = 0.020$ vs. the monthly-series 0.023) is reported in Appendix C.1.1.

Appendix B.2. Sector-level and country-level estimates

Table B.7 reports per-sector estimates for each of the 19 NAICS sectors in the BTOS data. The imitation rate q hits the lower bound of the optimisation in every sector and SE(q) dwarfs the point estimate, reflecting the well-documented difficulty of identifying imitation from a two-year window of an early-stage diffusion (Appendix C.1.2); the per-sector fit effectively reduces to the constant-innovation model $dF/dt = p(1 - F)$.

Table B.8 presents logistic diffusion estimates for 37 countries, combining OECD and Eurostat data sources.

Appendix B.3. International and cross-country patterns

Beyond the US time series used for estimation, cross-country data reveal both the dispersion in AI adoption and its national-level correlates. In Eurostat's 2025 cross-section of 33 European countries, individual generative-AI adoption ranges from 17.2% (Türkiye) to 56.3% (Norway), with Nordic and Baltic countries leading. Enterprise rankings reorder toward Denmark (42.0%), Finland (37.8%), and Sweden (35.0%)—reflecting country-specific factors beyond individual familiarity. Full country-level rates are available

Table B.7: Sector-level Bass parameter estimates (BTOS)

Sector	b_0	k	p	SE(p)	q	SE(q)	R^2
Construction	0.3922	0.1438	0.0016	0.0119	0.0010	0.8311	0.6636
Wholesale Trade	0.4384	0.1996	0.0026	0.0155	0.0010	0.6883	0.6191
Finance/Insurance	0.4373	0.0883	0.0071	0.0351	0.0010	0.6042	0.6180
Health Care	0.3272	0.0436	0.0050	0.0190	0.0010	0.4293	0.3822
Transportation	0.3703	0.1787	0.0014	0.0169	0.0010	1.3730	0.1952
Admin/Waste Services	0.3326	0.1273	0.0043	0.0249	0.0010	0.6878	0.1366
Professional Services	0.2972	0.0527	0.0126	0.0230	0.0010	0.2226	0.1303
Arts/Entertainment	0.2972	0.0877	0.0045	0.0309	0.0010	0.7486	0.1284
Education	0.2378	0.0094	0.0089	0.0567	0.0010	0.7545	-0.1040
Manufacturing	0.3762	0.0122	0.0023	0.0170	0.0010	0.8457	-0.1085
Real Estate	0.2361	-0.0123	0.0068	0.0319	0.0010	0.5511	-0.2254
Retail Trade	0.2832	0.0276	0.0029	0.0134	0.0010	0.5046	-0.2421
Other Services	0.2572	-0.1318	0.0022	0.0164	0.0010	0.8672	-0.3837
Agriculture	–	–	0.0025	1.1828	0.0010	47.7001	-0.3895
Accommodation/Food	0.4072	0.1628	0.0013	0.0108	0.0010	0.9486	-0.4181
Utilities	–	–	0.0043	0.3760	0.0010	9.5839	-0.5123
Information	0.2442	0.0294	0.0166	0.0771	0.0010	0.5806	-1.3027
Management	–	–	0.0149	1.0470	0.0010	9.9775	-2.9660
Mining	–	–	0.0030	0.4738	0.0010	21.8532	-21.8713

Notes: Two estimation procedures are applied independently per sector. Barrier parameters (b_0, k) are estimated by log-linear regression of the BTOS planned-versus-actual adoption gap on cumulative adoption; dashes denote sectors where this regression did not converge. Innovation and imitation rates (p, q) are estimated by nonlinear least-squares fits of $dF/dt = (p + qF)(1 - F)$, with standard errors from the Jacobian and R^2 measuring the fit of this specification.

from the cited Eurostat datasets (`isoc_ai_iaiu`, `isoc_eb_ai`); the distributional summary is reported in Panel E of Table B.6.

European AI adoption shows a steep age gradient: 63.8% of 16–24 year-olds have used generative AI versus 6.5% of those aged 65–74. The four-wave US Pew trajectory (13.6%, 17.7%, 22.7%, 33.4%) yields a weighted logistic fit of $K = 47.8\%$ ($R^2 = 0.998$), well below the W152 "any AI tool" rate of 57.1%, consistent with chatbots being one category within a broader adoption wave. The imitation coefficient q is indistinguishable from zero across all specifications; the truncation analysis in Appendix C.1.2 shows this pattern is expected from a two-year estimation window regardless of the true q.

Reasons for non-adoption. Eurostat also reports reasons why individuals did not use generative AI tools (dataset `isoc_ai_iaiuxr`). Across the EU27, the dominant reason is *no perceived need* (39.2%), followed by *did not know how to use* (8.4%), *did not know they existed* (5.2%), and *privacy/security*

Table B.8: Country-level logistic diffusion estimates

ISO2	Country	r	T_{50}	R^2	Latest (%)	Source
RS	Serbia	2.0566	2,024	0.9925	–	Eurostat
EL	Greece	0.8293	2,023	0.8072	–	Eurostat
KR	South Korea	0.7884	2,024	0.7907	30.2800	OECD
EE	Estonia	0.6674	2,027	0.9426	23.4026	OECD
LT	Lithuania	0.6031	2,027	0.7162	21.3034	OECD
CY	Cyprus	0.5212	2,025	0.9755	–	Eurostat
SE	Sweden	0.4894	2,026	0.9085	35.0402	OECD
AT	Austria	0.4752	2,027	0.9571	29.9459	OECD
BE	Belgium	0.4527	2,026	0.9665	34.5428	OECD
BA	Bosnia	0.4233	2,030	0.9748	–	Eurostat
NO	Norway	0.4183	2,027	0.8817	28.8939	OECD
NL	Netherlands	0.4037	2,027	0.9260	33.2079	OECD
LU	Luxembourg	0.4011	2,027	0.9353	33.6114	OECD
LV	Latvia	0.3909	2,030	0.9567	12.2139	OECD
CZ	Czechia	0.3861	2,029	0.8749	17.6049	OECD
FR	France	0.3569	2,030	0.8019	18.1592	OECD
HU	Hungary	0.3504	2,031	0.9073	10.3731	OECD
FI	Finland	0.3490	2,027	0.8569	37.8204	OECD
SK	Slovakia	0.3413	2,030	0.8681	17.9952	OECD
DE	Germany	0.3396	2,028	0.9363	25.9684	OECD
GR	Greece	0.3316	2,032	0.7605	8.9269	OECD
SI	Slovenia	0.3142	2,029	0.8239	21.6118	OECD
DK	Denmark	0.3102	2,027	0.6822	42.0324	OECD
ME	Montenegro	0.3037	2,032	0.9885	–	Eurostat
ES	Spain	0.2673	2,031	0.7916	20.2685	OECD
IT	Italy	0.2468	2,033	0.5194	16.4032	OECD
MT	Malta	0.2339	2,031	0.9709	–	Eurostat
PL	Poland	0.2228	2,036	0.7448	8.3624	OECD
HR	Croatia	0.2221	2,033	0.8516	15.1928	OECD
RO	Romania	0.2011	2,040	-0.1705	5.2108	OECD
BG	Bulgaria	0.1931	2,038	0.6524	8.5463	OECD
TR	Turkey	0.1717	2,040	0.1823	7.4125	OECD
PT	Portugal	0.1382	2,040	0.5502	11.5446	OECD
BR	Brazil	0.1107	2,040	-14,508	12.9400	OECD
IE	Ireland	0.1049	2,040	-0.0445	19.6416	OECD

Notes: r is the growth rate, T_{50} the projected year of 50% adoption, and "Latest" the most recent observed adoption rate; dashes mark Eurostat series with growth dynamics but no comparable level observation. Logistic curves are fitted with the ceiling fixed at $K = 1$, and countries are sorted by r.

concerns (4.4%). The dominance of "no need" as a barrier is consistent with the innovation-driven adoption pattern documented in this paper: the binding constraint is not awareness or skill but rather the perceived value proposition—exactly the mechanism captured by the innovation coefficient p_i rather than the imitation coefficient q_i in the Bass framework.

Turning from the cross-section to its determinants, we exploit the variation in our 45-country panel to examine what drives international differences in AI adoption: levels range from below 5% (Latin America) to above 40% (Denmark), with the highest rates concentrated in Northern Europe. Table B.9 reports OLS regressions of the latest AI adoption rate on potential determinants. The Going Digital composite score—capturing infrastructure, skills, innovation, and regulatory quality—is the strongest single predictor ($\hat{\beta} = 1.19$, $p < 0.001$, $R^2 = 0.31$); adding GDP per capita and R&D raises R^2 to 0.70, with R&D the most robust predictor ($\beta = 6.11$, $p < 0.01$). Concretely, raising Turkey's digital readiness to Denmark's level implies a roughly 16.5 pp increase in enterprise AI adoption. An EU dummy is insignificant.

Table B.9: Determinants of national AI adoption rates: cross-country OLS

Variable	(1)	(2)	(3)	(4)
Constant	−58.53***	−40.34	−13.65	−14.06
	(18.11)	(29.29)	(27.46)	(38.15)
Going Digital	1.191***	0.921**	0.276	0.314
	(0.294)	(0.460)	(0.483)	(0.488)
GDP/capita ($000s)		0.085	0.172**	0.163
		(0.081)	(0.070)	(0.113)
R&D (% GDP)			6.109***	6.073***
			(2.190)	(2.253)
EU member				−1.710
				(27.17)
R^2	0.313	0.454	0.698	0.699
Adj. R^2	0.292	0.396	0.647	0.628
N	35	22	22	22

Notes: The dependent variable is the latest enterprise AI adoption rate (%). Robust standard errors in parentheses. The sample narrows from 35 countries in Column 1, where the OECD Going Digital score is available, to 22 in Columns 2–4, which additionally require World Bank GDP per capita and R&D data; the Going Digital coefficient is stable across this restriction ($\hat{\beta} = 1.19 \rightarrow 0.92$). $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

As a check on external validity, the cross-country regression estimated on EU countries predicts non-EU adoption (South Korea, Japan, Turkey, Colombia) with $R^2 = 0.86$ and a mean absolute error of 3.0 pp, indicating the relationship generalises beyond Europe.

Appendix B.4. AI diffusion parameters in historical context

Table B.10 compares the Bass diffusion parameters estimated in this paper with canonical values for historical technologies. The substantially different estimation horizons (2 years for AI vs. 10–25 years for historical technologies) limit direct comparability of the q estimates: the truncation analysis in Appendix C.1.2 shows that historical technologies would also yield $\hat{q} \approx 0$ when estimated from only the first two years of data.

Table B.10: AI diffusion parameters versus historical technologies

Technology	Source	p	q	p/q	Window
Panel A: Historical benchmarks					
Average (Bass 1969)	Bass [3]	0.030	0.380	0.079	∼15 yr
Average (meta-analysis)	Mahajan et al. [53]	0.030	0.410	0.073	∼15 yr
Panel B: Consumer technologies					
Television	Bass [3]	0.028	0.251	0.112	∼25 yr
Mobile phones	Meade and Islam [56]	0.005	0.500	0.010	∼20 yr
Internet (households)	Meade and Islam [56]	0.020	0.400	0.050	∼15 yr
Panel C: Enterprise technologies					
E-commerce	Meade and Islam [56]	0.015	0.300	0.050	∼10 yr
Cloud computing‡	Meade and Islam [56]	0.015	0.350	0.043	∼10 yr
ERP systems	Meade and Islam [56]	0.008	0.250	0.032	∼15 yr
Robots (manufacturing)	Meade and Islam [56]	0.005	0.200	0.025	∼20 yr
Panel D: AI — This paper					
AI (individual)	This paper	0.273	$\approx 0^{\dagger}$	—	**2 yr**
AI (organisational)	This paper	0.066	$\approx 0^{\dagger}$	—	**2 yr**

Notes: All parameter values are annual rates; Window is the approximate estimation horizon. Panel D estimates are this paper's NLS results (Table 1), converted from monthly to annual; the AI horizon (∼2 yr) is far shorter than historical windows (10–25 yr), limiting direct comparability of q (Appendix C.1.2). †q is indistinguishable from zero for AI (F-test $p = 1.00$), so p/q is not reported. ‡Cloud computing parameters are representative estimates calibrated to the enterprise technology range in Meade and Islam [56]; dedicated long-run estimates are unavailable.

Appendix C. Robustness and sensitivity

Appendix C.1. Data and measurement validation

Appendix C.1.1. Validation of the monthly adoption series

The joint ODE estimation requires a monthly individual adoption path; four Pew waves alone cannot identify curve shape. This subsection validates the Google Trends–based monthly series (Appendix B.1) against three concerns: proportionality between search interest and adoption, keyword weighting sensitivity, and consistency with survey-only estimation.

Multi-anchor validation and joint ODE sensitivity. Single-anchor (W152-only) calibration overstates early adoption by up to 8.1 pp (Table C.11, Panel A), reflecting an early-2023 curiosity spike. Multi-anchor rescaling pins levels at all four Pew waves, eliminating the bias (MAE = 0.4 pp) and yielding $\hat{p}_i = 0.021$, $\hat{q}_i \approx 0$, $R^2 = 0.95$. To bound the concern that the within-interval GT shape reflects curiosity rather than adoption, we compare the joint ODE estimate under the GT-shape specification against a piecewise-linear interpolation that discards GT shape entirely: $\hat{\beta}$ moves from 0.010 to 0.011 ($< 10\%$), so Pew level anchors carry identification. Under multi-anchor calibration, all six keyword weighting schemes yield $\hat{p}_i \in [0.020, 0.021]$ (Panel B); Pew anchors pin levels at four dates, limiting variation to 0.001. Re-estimating the full coupled ODE under all three GT constructions confirms invariance: $\hat{\beta}_{\text{ODE}} \in [0.009, 0.011]$ (Table C.12, Panel B), with multi-anchor variants yielding three-fold smaller SEs and significance at $p < 0.004$.

Survey-only and joint ODE validation. We re-estimate (p_i, q_i) from survey data alone, compiling seven points: four Pew waves plus YouGov [57], Ipsos/KnowledgePanel [58], and CR/NORC AmeriSpeak [59]. Five specifications—Pew-only through Pew-excluded (3 non-Pew sources only)—all yield $\hat{p}_i \in [0.020, 0.024]$ with $\hat{q}_i = 0$ (Table C.12, Panel A). The Pew-excluded estimate ($\hat{p}_i = 0.024$, SE 0.008) independently corroborates the GT-based estimate of 0.023, ruling out circularity. The universality of $\hat{q}_i = 0$ is consistent

Table C.11: Google Trends proxy validation

Panel A: Single-anchor validation against Pew anchors

Wave (date)	Pew rate (%)	GT implied (%)	Deviation (pp)	±2 SE?
W123 (Mar 2023)	13.6	21.7	+8.1	No
W131 (Jul 2023)	17.7	20.4	+2.7	No
W142 (Feb 2024)	22.7	27.6	+4.9	No
W152 (Aug 2024)	33.4	33.4	+0.0	Yes

Panel B: Keyword weighting sensitivity

	Single anchor (W152)			Multi-anchor (4 waves)		
Weighting scheme	$\hat{p}_i$	$\hat{q}_i$	R^2	$\hat{p}_i$	$\hat{q}_i$	R^2
Baseline (manuscript)	0.0234	0.0010	0.338	0.0206	0.0010	0.952
ChatGPT only	0.0232	0.0010	0.261	0.0205	0.0010	0.941
Equal weights (all 7)	0.0287	0.0010	-2.307	0.0204	0.0010	0.883
Top-3 equal	0.0217	0.0010	0.882	0.0208	0.0010	0.913
Consumer-focused	0.0220	0.0010	0.414	0.0205	0.0010	0.948
Broad AI interest	0.0270	0.0010	0.220	0.0213	0.0010	0.893
Range (p_i)	[0.0217, 0.0287]			[0.0204, 0.0213]		
Pew-only (no GT)	—			0.0201	0.0000	0.788

Notes: Panel A: single-anchor calibration rescales the GT composite to match W152 (33.4%). Panel B: Bass estimates under alternative keyword weightings. Multi-anchor calibration interpolates between all four Pew waves and is therefore nearly invariant to keyword weights.

with innovation-driven diffusion, though Appendix C.1.2 cautions that a 2-year window may be too short to detect imitation. Panel B reports the full coupled ODE re-estimation under three individual-side specifications; $\hat{\beta}_{\text{ODE}}$ and boundary parameters are invariant to the GT construction, confirming that non-identification of p_o and α is structural to the within-country ODE: the two channels trade off along the likelihood ridge documented in Appendix C.2.1, and cross-country variation is the appropriate lever for identifying β (Section 4.5).

Appendix C.1.2. Truncation analysis of the $q \approx 0$ finding

The finding that $\hat{q} \approx 0$ for AI could reflect either a genuinely innovation-driven technology or a standard Bass process observed too early for the imitation component to become identifiable. We conduct a truncation experiment: synthetic monthly Bass data with known parameters for four technologies—Internet ($p = 0.020$, $q = 0.400$), mobile phones ($p = 0.005$, $q = 0.500$), e-commerce ($p = 0.015$, $q = 0.300$), and an AI-like stress test ($p = 0.020$, $q = 0.100$)—are generated with Gaussian noise ($\sigma = 0.01$, calibrated to Pew residuals), then re-estimated using only the first T years.

At $T = 2$ (matching our estimation window), a substantial share of 200 Monte Carlo replications fail to detect q: 63% for mobile phones, 49% for the AI-like case, 31% for e-commerce, and 11% for the Internet (Table C.13). By $T = 5$, the imitation coefficient is reliably recovered across all four cases. The implication: $\hat{q} = 0$ is *necessary* but not *sufficient* for concluding that AI lacks imitation dynamics—the finding is consistent with both genuinely innovation-driven adoption and standard Bass dynamics observed too early. Resolving this ambiguity requires re-estimation with longer time series.

Appendix C.2. Estimation robustness

Appendix C.2.1. Joint ODE bootstrap and profile likelihood

The joint ODE estimation in Section 4.2 yields boundary estimates for p_o and α (both ≈ 0), where asymptotic inference may not be reliable. We

Table C.12: Survey-only estimation and joint ODE sensitivity

Panel A: Individual adoption estimates across data specifications					
Specification	N	$\hat{p}_i$	$\hat{q}_i$	R^2	$\hat{K}$ (logistic)
A: Pew-only (4 waves)	4	0.0201	0.0000	0.788	1.000
		(0.0013)	(0.0081)		
B: Pew + prob. suppl. (6 pts)	6	0.0202	0.0000	0.672	1.000
		(0.0013)	(0.0079)		
C: Full series (7 pts)	7	0.0202	0.0000	0.811	1.000
		(0.0012)	(0.0078)		
D: Full series, equal weights	7	0.0205	0.0000	0.812	0.403
		(0.0020)	(0.0135)		
E: Non-Pew only (3 pts)	3	0.0244	0.0000	0.900	0.234
		(0.0079)	(0.0878)		
Prior estimates (for comparison):					
GT single-anchor (Table 2)	22	0.0220	0.0100	0.281	—
GT multi-anchor	22	0.0189	0.0010	0.906	—
Pew-only, GT-free	4	0.0201	0.0000	0.788	—
Joint ODE (Table 2)	76	0.0228	0.0000	0.420	—

Panel B: Joint ODE sensitivity to Google Trends calibration			
	Single-anchor (W152 only)	Multi-anchor (PW-linear)	Multi-anchor (GT-shape)
$\hat{p}_i$	0.023 (0.013)	0.017*** (0.003)	0.018*** (0.003)
$\hat{p}_o$	0.000^a	0.000^a	0.000^a
$\hat{\alpha}$	0.000^a	0.000^a	0.000^a
$\hat{\beta}$	0.009 (0.010)	0.011*** (0.004)	0.010*** (0.003)
R_i^2	0.42	0.95	0.97
R_o^2	0.91	0.91	0.91
N_i	22	21	21

Notes: In Panel A, time is measured in months since November 2022 and standard errors are in parentheses; specifications A–C use inverse-variance weighting, D uses equal weights, and E excludes all Pew waves. Panel B reports the restricted model ($q_i = q_o = 0$) with $b_0 = 0.329$ and $k = 0.068$ fixed. ^{a}At lower bound. *** $p < 0.01$.

Table C.13: Bass parameter recovery under data truncation

		Estimation horizon (years)					
Technology	True	2	3	5	10	15	30
Internet	$p = 0.020$	0.021	0.020	0.020	0.020	0.020	0.020
	$q = 0.400$	0.358	0.401	0.399	0.400	0.400	0.399
	$\%\ \hat{q} \approx 0$	*11%*	*0%*	*0%*	*0%*	*0%*	*0%*
Mobile phones	$p = 0.005$	0.007	0.007	0.005	0.005	0.005	0.005
	$q = 0.500$	$0.000^{\dagger}$	0.325	0.486	0.500	0.500	0.499
	$\%\ \hat{q} \approx 0$	*63%*	*16%*	*0%*	*0%*	*0%*	*0%*
E-commerce	$p = 0.015$	0.016	0.016	0.015	0.015	0.015	0.015
	$q = 0.300$	0.199	0.267	0.294	0.301	0.300	0.300
	$\%\ \hat{q} \approx 0$	*31%*	*0%*	*0%*	*0%*	*0%*	*0%*
AI-like (low-q)	$p = 0.020$	0.020	0.021	0.020	0.020	0.020	0.020
	$q = 0.100$	0.016	0.079	0.097	0.100	0.100	0.100
	$\%\ \hat{q} \approx 0$	*49%*	*30%*	*1%*	*0%*	*0%*	*0%*

Notes: Synthetic monthly Bass adoption data generated with known parameters, truncated to first T years, and re-estimated by nonlinear least squares ($p \in [0, 0.5]$, $q \in [0, 1]$). Gaussian noise $\sigma = 0.01$ (1 pp). Values are medians across 200 Monte Carlo replications. †Median $\hat{q} < 0.01$. $\%\ \hat{q} \approx 0$ = fraction of replications with $\hat{q} < 0.01$. At 2 years, a substantial fraction of replications fail to detect q, with the effect strongest for slower-growing technologies.

address this with bootstrap, profile likelihood, and alternative data specifications (Table C.14).

Residual bootstrap and diagnostic. We construct 1,000 bootstrap samples by resampling residuals and re-estimating the restricted model ($q = 0$). The bootstrap 95% CI for p_i is $[0.003, 0.028]$ (excludes zero); for β, $[0.004, 0.010]$ (excludes zero but is asymmetric, reflecting the flat profile above). The boundary parameters p_o and α hit their lower bounds in 70% and 67% of replications, confirming the boundary is a genuine likelihood feature. However, Ljung–Box tests reject i.i.d. residuals (individual $\hat{\rho}_1 = 0.85$, $p < 0.001$; organisational $\hat{\rho}_1 = 0.40$, $p = 0.002$), so bootstrap SEs understate uncertainty. We therefore anchor inference on the asymptotic SEs and profile likelihood, both agnostic to residual dependence: the asymptotic SE for β (0.010) exceeds the bootstrap SE (0.001) by an order of magnitude.

Profile likelihood. The 95% profile CI for β is $[0.000, 0.010]$, which includes zero: the joint ODE does not identify $\beta > 0$ from within-country time series alone (Figure C.6). This is a design feature—identification of $\beta > 0$ rests on the cross-country evidence in Section 4.5 ($\hat{\beta}_{\text{cross}} = 0.73$, $p < 0.001$). The within-country flatness reflects limited identifying power: 22 monthly observations of a single rising F_i leave βF_i and the constant organisational hazard near-collinear. The cross-country regression circumvents this with 26 independent country-level (F_i, F_o) pairs at the same calendar date. The simulation's $\beta_{max,h}$ are therefore calibrated to the cross-country estimates (Section 5.1; Appendix C.3.1), not to the joint-ODE point estimate. The profile for p_i is better identified ($[0.019, 0.024]$), consistent with the survey-only estimates in Appendix C.1.1. The effective dimensionality of the individual series is $\sim$5 (four Pew anchors plus between-anchor shape), not the nominal $N_i = 22$; the Jacobian eigenvalue drop from 0.50 to 0.0006 confirms that α and p_o lie in the near-null space.

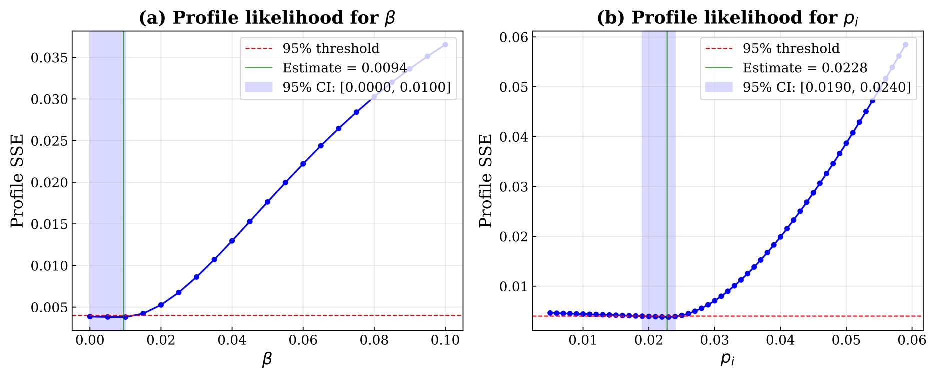

Fig. C.6. Profile likelihood for β and p_i in the within-country joint ODE.

Extended individual series. Re-estimating with Pew-only (4 waves) and Pew + RPS (9 points) yields stable $\hat{p}_i$ (range $[0.020, 0.023]$) but $\hat{\beta}$ moves toward zero under the extended specification (Panel B of Table C.14), consistent with the flat profile and limited within-country information about β.

Appendix C.2.2. Sensitivity to first-stage barrier estimates

The two-stage estimation fixes the barrier parameters b_0 and k at their BTOS first-stage values before estimating the remaining ODE parameters; if those first-stage estimates are biased, the second-stage parameters inherit the bias. We probe sensitivity within the exponential family $b_t = b_0\, e^{-k(F_i+F_o)}$, which is not a free functional-form choice but the parametric model under which the BTOS planned-vs-actual gap is estimable by log-linear regression (Appendix B.1, Appendix B); alternative functional families (logistic, power-law) would require a different estimator with its own identifying restrictions and are not directly recoverable from the BTOS data. We re-estimate the joint ODE across a 3×3 grid: $b_0 \in \{0.311, 0.329, 0.347\}$ and $k \in \{0.046, 0.068, 0.090\}$ (± 1 SE around each baseline).

Across all 9 perturbations, $\hat{p}_i$ is completely invariant at 0.0228, the boundary parameters $\hat{p}_o$ and $\hat{\alpha}$ remain at their lower bounds, and the fit statistics $R_i^2 = 0.42$ and $R_o^2 = 0.913$ are unchanged. The invariance of $\hat{p}_i$ is mechanical, not informative: (b_0, k) enter only the organisational equation, and because

Table C.14: Joint ODE inference: bootstrap CIs and extended-series robustness

Panel A: Bootstrap and profile likelihood confidence intervals

		Asymptotic		Bootstrap		Bnd. hit
	Estimate	SE	95% CI	SE	95% CI	(%)
p_i	0.0228	0.0126	[0.0010, 0.0475]	0.0074	[0.0027, 0.0278]	1.0
p_o	0.0001	0.0026	[0.0001, 0.0052]	0.0004	[0.0001, 0.0014]	70.0
α	0.0000	0.3052	[0.0000, 0.5000]	0.1753	[0.0000, 0.5000]	66.5
β	0.0094	0.0101	[0.0000, 0.0291]	0.0013	[0.0038, 0.0095]	0.4

Panel B: Joint ODE estimates under alternative individual series

	Original GT (22 months)	Pew-only (4 waves)	Extended (Pew + RPS)
$\hat{p}_i$	0.0228*	0.0010	0.0010
	(0.0126)	(0.0043)	(0.0027)
$\hat{p}_o$	0.0001	0.0001	0.0022
	(0.0026)	(0.0024)	(0.0020)
$\hat{\alpha}$	0.0000^{a}	0.1660*	0.1723***
	(0.3052)	(0.0910)	(0.0448)
$\hat{\beta}$	0.0094	0.0115	0.0015
	(0.0101)	(0.0103)	(0.0083)
R_i^2	0.42	0.86	0.94
R_o^2	0.91	0.84	0.77
N_i	22	4	9

Notes: Panel A reports asymptotic and bootstrap inference for the restricted joint ODE. Asymptotic standard errors are from the linearised Jacobian; bootstrap standard errors and confidence intervals are from a residual bootstrap (B = 1,000); profile-likelihood intervals are shown in Figure C.6. "Bnd. hit" is the share of bootstrap replications in which the estimate reached a parameter bound. Ljung–Box tests reject i.i.d. residuals ($\hat{\rho}_1 = 0.85$ individual, 0.40 organisational), so the bootstrap standard errors understate uncertainty. Panel B re-estimates the restricted model ($q_i = q_o = 0$; $b_0 = 0.329$, $k = 0.068$ fixed; $N_o = 54$) under three individual-adoption series. ^{a}At lower bound. $^{***}p < 0.01$, $^{*}p < 0.10$.

$\hat{\alpha} \approx 0$ in the point estimate the individual ODE has no functional dependence on the barrier parameters, so $\hat{p}_i$ is guaranteed to be invariant by the model structure. The informative quantity is $\hat{\beta}$, which varies within a 6% band: $\hat{\beta} \in [0.0091, 0.0097]$, CV $= 2.2\%$, with higher b_0 requiring slightly larger β to fit the same organisational trajectory and higher k allowing slightly smaller β as barrier erosion compensates. The qualitative conclusions are therefore insensitive to first-stage barrier uncertainty for the parameter that the barrier perturbations could plausibly affect.

Appendix C.2.3. Cross-group regression robustness

The cross-country OLS estimate $\hat{\beta}_{\text{cross}} = 0.73$ ($p < 0.001$, Table 4) could reflect shared digital readiness rather than a genuine cross-group channel [60]. We address this concern through a reverse-direction regression, three alternative identification strategies, and finite-sample inference validation.

Reverse regression. The reverse regression—individual GenAI adoption on enterprise AI adoption—yields $\hat{\alpha}_{\text{cross}} = 0.69$ ($p < 0.001$), comparable to the forward $\hat{\beta}_{\text{cross}} = 0.73$ (ratio 1.06), precisely the near-symmetry expected under the reflection problem. Controlling for Going Digital attenuates $\hat{\alpha}_{\text{cross}}$ by 33% vs. 48% for $\hat{\beta}_{\text{cross}}$; this 15 pp gap motivates the asymmetric scaling factors ($\times 0.32$ for α vs. $\times 0.25$ for β; Section 5.1). Table C.15 reports the full results.

Alternative identification. Three strategies probe the cross-group association beyond OLS (Table C.16).

Panel fixed effects. Eurostat enterprise panel data (27 countries, 2021–2025) with country fixed effects yields $\hat{\beta}_{\text{panel}} = 0.24$ ($p = 0.002$), identified from within-country changes. First-differences yield a consistent $\hat{\beta}_{\text{FD}} = 0.17$ ($p = 0.075$).

Instrumental variables. Instrumenting individual GenAI adoption with demographic composition (first-stage $F = 30.9$) yields $\hat{\beta}_{\text{IV}} = 0.66$ ($p <$

Table C.15: Reverse cross-group influence: enterprise AI and individual GenAI adoption

	(R1)	(R2)	Comparison
DV:	Individual GenAI (%)		Forward (Table 4)
Enterprise AI (%)	0.688***	0.462*	
	(0.136)	(0.206)	
Going Digital score		0.309	
		(0.488)	
Forward coefficient			
Individual GenAI (%)			0.728*** / 0.378*
R^2	0.501	0.424	0.501 / 0.600
Countries	26	23	26 / 23
Attenuation by GD control	33%		48%
$\hat{\beta}/\hat{\alpha}$ ratio	—		1.06 / 0.82

Notes: Robust SEs in parentheses. Column (R1): bivariate regression of individual GenAI adoption on enterprise AI adoption. Column (R2): controlled for OECD Going Digital score. Comparison column reproduces the corresponding forward-direction coefficients from Table 4 (Columns 1 and 3). Attenuation = percentage reduction in the slope coefficient when the Going Digital control is added. The near-symmetric bivariate coefficients ($\hat{\beta}/\hat{\alpha} = 1.06$) are consistent with the reflection problem [60]. $^{***}p < 0.001$, $^{*}p < 0.05$.

Table C.16: Cross-group influence: panel estimation of β

Specification	$\hat{\beta}$	SE	p-value	R^2	N
Cross-section (2025)	0.598***	0.213	0.010	0.247	26
Pooled OLS (all years)	0.355***	0.094	0.000	0.122	105
Country FE + year FE	0.058	0.043	0.180	0.797	105
Country FE × latest year	0.243***	0.076	0.002	0.797	105
First-differences	0.174*	0.093	0.075	0.137	24
SME only (2025)	0.589***	0.209	0.010	0.248	26
SME only (2025)	0.960***	0.309	0.005	0.279	27

Notes: Dependent variable: enterprise AI adoption rate (%). Individual AI adoption measured from 2025 OECD/Eurostat data. Country FE specification uses ind_ai × year interactions to identify differential enterprise growth associated with individual adoption levels. First-differences use ΔF_o between consecutive years. Robust SEs. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

0.001), 9% below OLS. A Hausman test ($p = 0.28$) is consistent with modest confounding but cannot reject OLS consistency. The most plausible exclusion-restriction violation—the workforce-composition channel—is partially addressed by the controlled OLS specifications (Going Digital, GDP, R&D). We report the IV as triangulation rather than a clean causal estimate.

Lagged specification. Enterprise adoption *growth* (2024→2025) on individual adoption levels yields $\hat{\beta}_{\text{lag}} = 0.046$ ($p = 0.008$), with a larger coefficient for large firms (0.17, $p < 0.001$). The effect is absorbed when Going Digital is added ($p = 0.48$), indicating the channel operates through shared digital infrastructure rather than direct transmission.

All four strategies confirm a positive cross-group association; the magnitude attenuates as confounders are absorbed, consistent with a genuine but partially confounded effect. Sensitivity of the simulation's β_{max} across this range is tested in Appendix C.3.1.

Finite-sample inference. The cross-country regression uses only $N = 23$–26 countries. Wild bootstrap with Rademacher weights [61, 62] yields $p < 0.001$ for both bivariate and Going-Digital-controlled specifications; randomisation inference (10,000 permutations) yields $p < 0.0001$ for both. The triple agreement (asymptotic, wild bootstrap, randomisation) confirms that significance is not a small-sample artefact.

Appendix C.3. Calibration robustness

Appendix C.3.1. Scaling-factor sensitivity

The reduced-form estimates of Section 4 are converted into ODE flow rates by empirical scaling factors—a uniform $\times 0.25$ for the cross-group parameters $\beta_{max,h}$ and $\times 0.20/\times 0.15$ for the individual and organisational innovation coefficients. The conversion is structural: regression and logistic estimates are stock-on-stock or interval-growth objects, while the ODE parameters are per-period flow rates, so translation requires dividing by the operating horizon and accounting for the $(1 - F)$ saturation factor that the reduced forms omit. We report sensitivity to both sets of factors in turn.

Cross-group scaling factor. The parameters $\beta_{max,h}$ derive from the firm-size-stratified OECD cross-country regression coefficients (Section 4.5; $\hat{\beta}_{\text{large}} = 1.17$, $\hat{\beta}_{\text{medium}} = 1.04$, $\hat{\beta}_{\text{small}} = 0.65$) by a uniform $\times 0.25$, yielding $\beta_{max,h} = [0.292, 0.261, 0.163]$ for large/medium/small firms. We report sensitivity across the range $[0.10, 0.50]$ and validate the baseline against a held-out 2025 Eurostat wave.

Out-of-sample validation. A natural concern is that the scaling factor was chosen to minimise MAE against the same data the model is supposed to forecast. To address this, we split the validation: the in-sample objective is combined MAE against all Pew waves (Jan 2023–Aug 2024) and pre-2025 Eurostat enterprise observations (2021, 2023, 2024); the 2025 Eurostat wave is held out as an out-of-sample target. Table C.17 reports the validation surface across the grid. The baseline ×0.25 minimises the OOS MAE (0.24 pp on the held-out 2025 wave); the in-sample MAE surface, by contrast, is flat—varying by only 0.36 pp across the full grid (2.44–2.80 pp)—so ×0.25 is within 0.06 pp of the in-sample minimum at ×0.30 (2.50 vs. 2.44 pp). The OOS column traces a sharp V-shape around the baseline: moving right, MAE rises to 1.49 pp (×0.30), 3.93 pp (×0.40), and 6.31 pp (×0.50); moving left, to 1.03 pp (×0.20), 2.31 pp (×0.15), and 3.62 pp (×0.10). Higher scaling factors fit the pre-2025 organisational series too aggressively and over-predict the 2025 wave, while lower factors under-fit it. The baseline therefore generalises forward by one full Eurostat wave, ruling out the interpretation that ×0.25 is an in-sample artefact.

In-sample sensitivity surface. Across the full grid, the individual adoption MAE is nearly invariant to the scaling factor (∼3.3–3.6 pp) because β enters only the organisational adoption equation. The pre-2025 organisational MAE traces a U-shaped curve with a flat minimum around ×0.25–×0.40, all within

Table C.17: Sensitivity of simulation results to the cross-group scaling factor

		Organisational adoption (%)				MAE (pp)		
Factor	β_{max} range	2025	2028	2030	2033	In-samp.	OOS 2025	BCR
×0.10	[0.07, 0.12]	15.6	30.7	41.3	55.9	2.66	3.62	0.22
×0.15	[0.10, 0.18]	16.9	35.0	47.4	63.2	2.60	2.31	0.23
×0.20	[0.13, 0.23]	18.2	39.1	52.8	69.2	2.55	1.03	0.24
×0.25 ←	[0.16, 0.29]	19.5	42.9	57.6	74.0	2.50	$0.24^{\ddagger}$	0.25
×0.30	[0.20, 0.35]	20.7	46.5	61.9	78.1	$2.44^{\dagger}$	1.49	0.26
×0.40	[0.26, 0.47]	23.2	53.0	69.1	84.2	2.58	3.93	0.27
×0.50	[0.33, 0.58]	25.6	58.7	74.9	88.5	2.80	6.31	0.27

Notes: Scaling factor maps annual cross-country $\hat{\beta}_{h,\text{cross}}$ to simulation $\beta_{max,h}$: $\beta_{max,h} = \hat{\beta}_{h,\text{cross}} \times$factor, applied uniformly across firm archetypes ($\hat{\beta}_{\text{large}} = 1.17$, $\hat{\beta}_{\text{medium}} = 1.04$, $\hat{\beta}_{\text{small}} = 0.65$). In-sample MAE = combined error against Pew (all waves) and pre-2025 Eurostat (2021, 2023, 2024); OOS 2025 = held-out error against the 2025 Eurostat wave. BCR = benefit-cost ratio for Education & Training policy at 5-year horizon. ← baseline specification (×0.25); † in-sample MAE minimiser; ‡ OOS MAE minimiser.

0.1 pp of each other, so in-sample fit alone barely distinguishes the candidate factors. This flatness is exactly why the OOS test is informative: when the in-sample surface cannot identify a unique best factor, the held-out wave provides the discriminating evidence.

Robustness of qualitative conclusions. Qualitative conclusions are robust across the full range: the functionally asymmetric influence structure (the β channel dominates organisational adoption despite $\beta_{max} \approx \alpha_{max}$), the ranking of policy interventions (Education & Training most cost-effective), and the firm-size gap pattern all hold at every scaling factor. Quantitative organisational adoption forecasts are more sensitive: projected 2030 adoption ranges from 41.3% (×0.10) to 74.9% (×0.50), a 34 pp span (Table C.17, columns 3–6). The policy benefit-cost ratio (BCR) for Education & Training varies from 0.22 to 0.27, preserving the qualitative conclusion but with material quantitative uncertainty.

Innovation scaling factors. The innovation coefficients $p_{i,g}$ and $p_{o,h}$ are derived from reduced-form logistic growth rates by applying scaling factors of 0.20 (individuals) and 0.15–0.19 (organisations); we provide a formal derivation and report sensitivity to these choices below.

Functional form mismatch. We estimate sector-level diffusion using the logistic functional form because the Bass model with $q \neq 0$ is poorly identified in 2-year windows (Appendix C.1.2: 11–63% of replications fail to recover q at $T = 2$ across the four synthetic technologies); after estimating r under the logistic form, we translate to the simulation's Bass-with-$q = 0$ specification via an empirical scaling factor. The logistic ($dF/dt = rF(1 - F)$) and Bass-with-$q = 0$ ($dF/dt = p(1-F)$) differ structurally: the logistic has an interior inflection at $F = 0.5$ while Bass-with-$q = 0$ decelerates monotonically from $t = 0$. Matching instantaneous growth at $t = 0$ gives $p = rF_0$, but matching the trajectories over the 2-year estimation window requires a larger p because the logistic accelerates over the interval while Bass does not. Concretely, for the mid individual archetype ($r = 0.80$, $F_{i,0} = 0.12$) the logistic trajectory reaches $F_i \approx 40.5\%$ at $t = 2$ while the instantaneous-match Bass($q = 0$) with $p = rF_{i,0}$ reaches only $\approx 27.4\%$—a 13 pp shortfall that the empirical scaling factor closes. Mean-squared-error trajectory matching against the 2-year logistic data yields scaling factors of 0.18–0.22 for individuals (mean 0.21) and 0.04–0.19 for organisations (mean 0.11, $F_{o,0} = 0.08$). The calibrated baseline values 0.20 and 0.15–0.19 sit within these analytical ranges. The organisational factors (0.15 for medium and small firms, 0.19 for large) sit above the trajectory-matched mean (0.11) because in the coupled ODE the term βF_i provides additional growth that substitutes for p_o, lowering the autonomous innovation required to match the observed trajectory; large firms, whose autonomous adoption capacity is greatest, are calibrated at the upper end of the trajectory-matched range to reproduce their faster observed T_{50}. Sensitivity in Appendix C.3.1 confirms robustness across $[0.10, 0.30]$.

Asymmetry. The ratio of individual to organisational scaling factors ($0.20/0.15 = 1.33$) is consistent with the trajectory-matched ratio ($0.21/0.11 \approx 1.9$) and reflects the stronger functional role of cross-group support for organisations. Although $\beta_{max} \approx \alpha_{max}$ after symmetric scaling, organisations require less autonomous innovation (p_o) to match observed adoption because $\beta F_{i,\text{total}}$ contributes substantially to their adoption rate, while the α channel is redundant for individuals whose high $p_{i,g}$ already drives rapid uptake. The β–p_o trade-off is observable empirically: across the β-scaling sensitivity grid (Appendix C.3.1), in-sample MAE is approximately flat across β-scaling factors $[0.20, 0.40]$ when p_o-scaling is held at baseline, indicating that the organisational trajectory can be fit by trading off β_{max} against p_o. The asymmetric ratio is the centring choice within this flat region of the loss surface, not a sharp identification; the joint scaling grid in Appendix C.3.1 probes this trade-off directly.

Sensitivity analysis. Table C.18 reports how simulation results change as each scaling factor is varied while holding the other at its baseline value. For the individual factor (Panel A), validation MAE traces a U-shaped curve with a minimum at $\times 0.12$ (1.27 pp), close to the baseline $\times 0.20$ (2.17 pp). Individual adoption at 2030 ranges from 68.7% ($\times 0.10$) to 90.4% ($\times 0.30$), a 22 pp span. For the organisational factor (Panel B), the baseline $\times 0.15$ minimises combined MAE (2.17 pp), and organisational adoption at 2030 ranges from 50.7% ($\times 0.05$) to 62.8% ($\times 0.25$), a 12 pp span. The Education & Training policy BCR is stable across all specifications (0.22–0.25), preserving the qualitative ranking of policy interventions.

Joint variation. The preceding one-at-a-time analyses cannot capture the β–p_o substitutability noted above. Table C.19 reports F_o at 2030 across a 3×3 grid: β-scaling factors 0.15, 0.25, and 0.35, crossed with p_o-scaling factors 0.10, 0.15, and 0.20. The baseline cell ($\times 0.25$, $\times 0.15$) coincides with the joint MAE minimiser.

Four qualitative conclusions are universally robust across all nine cells: (i) individuals always lead organisations in adoption timing ($F_i > F_o$ at every horizon); (ii) adoption always drives productivity above the baseline ($P_{\text{final}} \in [1.41, 1.53]$); (iii) the firm-size adoption gap is always positive and

Table C.18: Sensitivity of simulation results to innovation coefficient scaling factors

		Adoption at 2030 (%)				
Factor	p range[b]	F_i	F_o	ΔF[a]	MAE (pp)	BCR
Panel A: Individual scaling factor varied[c]						
×0.10	[0.064, 0.102]	68.7	51.7	-14.2	1.72	0.22
×0.12	[0.076, 0.122]	72.3	53.2	-10.6	1.27	0.23
×0.15	[0.095, 0.153]	77.0	55.0	-5.9	1.34	0.24
×0.20 ←	[0.127, 0.204]	82.9	57.6	+0.0	2.17	0.25
×0.25	[0.159, 0.255]	87.2	59.6	+4.3	3.72	0.25
×0.30	[0.191, 0.306]	90.4	61.3	+7.5	5.20	0.24
Panel B: Organisational scaling factor varied[d]						
×0.05	[0.013, 0.016]	81.5	50.7	-6.5	2.78	0.25
×0.08	[0.020, 0.026]	81.9	52.8	-4.4	2.59	0.25
×0.10	[0.025, 0.033]	82.2	54.1	-3.1	2.47	0.25
×0.15 ←	[0.038, 0.049]	82.8	57.2	+0.0	2.17	0.25
×0.20	[0.051, 0.065]	83.4	60.1	+2.9	2.29	0.25
×0.25	[0.064, 0.082]	83.9	62.8	+5.5	2.64	0.25

[a] Deviation from baseline 2030 adoption in the varied margin: ΔF_i in Panel A, ΔF_o in Panel B.
[b] Range of the *varied* innovation coefficient across the three heterogeneity groups.
[c] Organisational scaling factor fixed at ×0.15, with $p_o \in [0.038, 0.049]$.
[d] Individual scaling factor fixed at ×0.20, with $p_i \in [0.127, 0.204]$.
Notes: Scaling factor maps reduced-form logistic growth rate r to Bass innovation coefficient: $p = r \times \text{factor}$. MAE = mean absolute error against Pew (individual) and Eurostat (enterprise) observed data. BCR = benefit-cost ratio for Education & Training policy at 5-year horizon. Arrow marks the baseline specification.

Table C.19: Joint scaling factor robustness: organisational adoption F_o at 2030

	p_o-scaling		
β-scaling	×0.10	×0.15	×0.20
×0.15	42.9%	46.8%	50.5%
×**0.25** (baseline)	54.1%	**57.2%**	60.1%
×0.35	63.0%	65.5%	67.8%

Notes: Each cell reports aggregate organisational adoption F_o at 2030 under the corresponding (β-scaling, p_o-scaling) pair, holding individual scaling at ×0.20 and barriers/productivity at baseline. Bold cell = baseline; this cell coincides with the joint MAE-minimising candidate within the grid: the baseline β-scaling (×0.25) minimises out-of-sample MAE in Appendix C.3.1, and the baseline p_o-scaling (×0.15) minimises one-at-a-time MAE in Appendix C.3.1. Other outcomes (suppressed for compactness) range over the grid: $T^o_{50} \in [4.8, 8.3]$ years, large-minus-small firm gap at Year 5 $\in [29.8, 37.8]$ pp, and F_i at 2028 $\in [67.7\%, 71.5\%]$.

substantial (30–38 pp at Year 5); and (iv) individual adoption is highly stable (F_i at 2028 varies by only 3.8 pp across the entire grid, from 67.7% to 71.5%). Quantitative organisational adoption forecasts vary meaningfully—F_o at 2030 spans 42.9% to 67.8%, a 25 pp range—but the qualitative ordering and policy-relevant structure are preserved. The most conservative combination (β-scaling = 0.15, p_o-scaling = 0.10) yields organisational T_{50} in 2031.3, while the most aggressive ($\times$0.35, $\times$0.20) yields 2027.8.

Appendix C.3.2. Sensitivity to a permanent barrier floor

The baseline barrier specification $b_t = b_0\, e^{-k(F_i+F_o)}$ decays monotonically to zero, implying that all organisational frictions are transitory. However, the BTOS planned–actual gap evidence (Section 4.2) suggests otherwise: planned adoption grows 1.3 times faster than actual adoption, and approximately 26% of firms that report planning to adopt AI never follow through. This persistent leakage motivates a modified barrier function with a permanent floor:

$$b_t = b_\infty + (b_0 - b_\infty)\, e^{-k(F_i+F_o)}, \quad b_\infty \geq 0, \tag{C.1}$$

where b_∞ represents the fraction of organisational barriers that are permanent (regulatory constraints, structural incompatibility, or persistent managerial resistance) rather than transitory. The empirically motivated value is $b_\infty \approx 0.08$, computed as $0.26 \times \bar{b}_0 = 0.26 \times 0.33$.

Table C.20 reports simulation results across $b_\infty \in [0, 0.15]$, holding all other parameters at their baseline values. The results are insensitive to the barrier floor: even at $b_\infty = 0.15$ (nearly double the empirically motivated value), organisational adoption at year 10 drops by only 0.6 pp (from 66.8% to 66.2%), T_{50} increases by 0.06 years, and the productivity index falls by less than 0.01. The firm-size gap is similarly unaffected. The insensitivity arises because the barrier enters as a multiplicative factor $(1 - b_t)$ on the hazard rate; a permanent floor of $b_\infty = 0.15$ reduces the long-run hazard by 15%, but does not alter the steady state (which remains at $F = 1$ for any positive innovation rate). The floor slows convergence but does not create a low-adoption trap, because the innovation coefficients $p_{o,h}$ keep the adoption hazard strictly positive regardless of the barrier level.

This is a deliberate feature of our specification rather than a robustness finding. The paper's identification target is the next-decade diffusion trajectory, not the long-run absorbing state: we report 2025–2033 forecasts and policy comparisons over a 10-year horizon, where the steady state is far outside the analysis window. A model that admits genuine long-run incomplete adoption—i.e., a stable interior equilibrium $F_o^\star < 1$—would require either heterogeneous innovation rates with positive mass at $p_{o,h} = 0$ (a subset of firms that never adopt regardless of barriers) or a barrier that enters additively in the hazard rather than multiplicatively (so that $b_\infty > 0$ shuts down the hazard entirely once F is high enough). Both specifications are reasonable for analyses centred on the steady state—e.g., long-run welfare comparisons or absorbing-fraction estimation—but neither is needed for next-decade forecasts where F_o remains well below 1 throughout. The barrier-floor robustness here documents that our calibration choices for the within-horizon dynamics are not contaminated by an implicit assumption about the long-run state. All qualitative conclusions within the analysis horizon—policy rankings, the heterogeneity premium, and the firm-size gap—are fully robust to this specification change.

Table C.20: Sensitivity to a permanent barrier floor (b_∞)

b_∞	F_o (Yr 5)	F_o (Yr 7)	F_o (Yr 10)	T_{50}^o (yr)	Firm gap (pp)	P_{10}
0.00	**38.9%**	**51.5%**	**66.8%**	**6.8**	**35.6**	**1.456**
0.05	38.8%	51.4%	66.6%	6.8	35.7	1.455
0.08	38.7%	51.3%	66.5%	6.8	35.7	1.455
0.10	38.7%	51.3%	66.4%	6.8	35.7	1.454
0.15	38.6%	51.1%	66.2%	6.8	35.7	1.454

Notes: Bold row = baseline ($b_\infty = 0$). b_∞ is the permanent barrier floor: $b_t = b_\infty + (b_0 - b_\infty)\, e^{-k(F_i+F_o)}$. The value $b_\infty = 0.08$ corresponds to the BTOS finding that ≈26% of firms planning AI adoption never follow through ($0.08 \approx 0.26 \times \bar{b}_0$). F_o = aggregate organisational adoption. Firm gap = large-minus-small firm adoption at Year 5. P_{10} = productivity index at Year 10. All other parameters held at baseline values (Table D.24).

Appendix C.3.3. Saturating cross-group influence

The baseline simulations use the linear specialisation of the cross-group influence ($\beta_h F_i = \beta_{max,h} F_i$, $\alpha_g F_o = \alpha_{max,g} F_o$), because the cross-country evidence of Section 4.5 identifies a linear individual–enterprise slope and observed adoption lies below the range over which the saturating form (Equations 7–8) departs materially from linearity. Here we verify that adopting the saturating form instead leaves the paper's conclusions intact. We set $\gamma_{i,h} = \gamma_{o,g} = 1$ and re-calibrate the cross-group scaling factor to the same out-of-sample target used in Appendix C.3.1; because the saturating transform discounts the influence term, matching the held-out 2025 enterprise anchor raises the factor from $\times 0.25$ to $\times 0.275$.

Table C.21 compares the two specifications. Every qualitative conclusion is preserved. The cross-group channel remains the dominant accelerator of organisational adoption: shutting it down ($\alpha = \beta = 0$) leaves organisational adoption at 36.0% at the 2033 horizon under both specifications, identical because the functional form of the cross-group term is irrelevant once it is switched off. The firm-size adoption gap at year 5 is essentially unchanged (36.4 versus 37.4 pp). The headline organisational forecast is modestly lower under saturation (68.8% versus 74.1% at 2033, with 50% reached in mid-2029 rather than late 2028), reflecting the saturating discount at high individual adoption. The within-population T_{50} dispersion across firm-size archetypes, the object heterogeneity is introduced to capture, is if anything sharper under saturation (6.03 versus 5.15 years), bracketing the six-year empirical spread of Table 3. Re-running the three interventions under the saturating calibration confirms the policy ranking is unaffected: at the five-year horizon Education and Training remains the most effective per unit of intensity (effectiveness ratio 0.24, versus 0.25 in the linear baseline), the Education > Subsidies > R&D ordering is preserved, and Direct Subsidies remains the most size-equalising intervention (narrowing the 2033 Large−Small adoption gap by 7.4 pp, versus 7.8 pp under the linear baseline).

Table C.21: Robustness to saturating cross-group influence

	Linear baseline	Saturating (re-calibrated)
Cross-group scaling factor	$\times 0.25$	$\times 0.275$
F_o at 2029 (%)	50.6	47.2
F_o at 2033 (%)	74.1	68.8
F_i at 2033 (%)	93.6	91.9
Organisational T_{50} (year)	2028.9	2029.5
T_{50} spread, Large–Small (yr)	5.15	6.03
Large$-$Small gap at year 5 (pp)	37.4	36.4
No cross-group: F_o at 2033 (%)	36.0	36.0
Education effectiveness ratio, 5-yr	0.25	0.24
Subsidies: 2033 gap narrowing (pp)	7.8	7.4

Notes: Both columns use the calibrated heterogeneous parameters of Table D.24. The saturating column sets $\gamma_{i,h} = \gamma_{o,g} = 1$ in Equations 7–8 and re-calibrates the cross-group scaling factor to the held-out 2025 Eurostat enterprise anchor (Appendix C.3.1). The empirical T_{50} spread across firm sizes is six years (Table 3). The effectiveness ratio is adoption gain (pp) per unit of intervention intensity, as in Table 5.

Appendix C.4. Core-result summary

Table C.22 consolidates the three robustness checks referenced in Section 4.6: the logistic-vs-Gompertz comparison for S-curve shape, outlier sensitivity for the Going Digital–adoption regression, and unweighted-vs-weighted Pew adoption rates. Each check's qualitative conclusion is stated in the main text; the table reports the underlying estimates. Two specific patterns warrant comment. First, the Gompertz growth rate ($r = 0.109$) is not directly comparable to the logistic growth rate ($r = 0.269$) because the two functional forms parameterise time differently: the logistic has r multiplying linear time in the hazard $rF(1-F)$, while the Gompertz has r multiplying $\ln(1/F)$ in $rF\ln(1/F)$. The comparable quantity is the fit ($R^2 = 0.813$ vs. 0.799, a difference of 1.4 percentage points) and the implied trajectory shape, both of which are nearly identical across the two specifications. Second, the outlier-treatment row drops one observation from each tail of the Going Digital distribution (the highest and lowest scoring countries), reducing N from 27 to 25 and the slope from 1.191 to 1.010—a 15% movement that reflects normal sensitivity of a cross-country OLS slope to support reduction, not a fragile finding. The slope remains statistically significant ($p < 0.001$) and the qualitative conclusion (Going Digital is the strongest single predictor of adoption) is unchanged.

Table C.22: Robustness of S-curve fit, cross-country regression, and survey-weighted adoption

Finding tested	Category	Specification	Est.	R^2
Adoption curve shape (S-curve robustness)				
S-curve	Functional form	Logistic (baseline)	0.269	0.813
S-curve	Functional form	Gompertz	0.109	0.799
Cross-country regression (Going Digital–adoption)				
Going Digital	Outlier treatment	Baseline	1.191	0.313
Going Digital	Outlier treatment	Winsorize 5/95	1.142	0.315
Going Digital	Outlier treatment	Drop top/bottom	1.010	0.395
Survey methodology (adoption level)				
Adoption rate	Survey weights	Unweighted (W152)	0.603	—
Adoption rate	Survey weights	Weighted (W152)	0.571	—

Notes: The Estimate column reports the growth rate r for the functional-form rows, the regression coefficient for the outlier-treatment rows, and the adoption rate for the survey-weights rows. The survey-weights row reports the broader "used any AI tool" variable (57.1% weighted), whereas our estimation uses the narrower "AI chatbot" variable (33.4% weighted); the 3.2 pp weighting effect is comparable across both variables.

Appendix D. Simulation calibration and robustness

Appendix D.1. Extended model parameterisation

Table D.24 (reported with the comparator table at the end of this section) lists the calibrated values used in the extended model and their empirical anchors. Three rows merit comment. First, $P_{max} = 2.0$ and $P_{initial} = 1.0$ are *normalisations*: the model's productivity index is dimensionless, so its scale and origin are set without loss of generality, and ratios $P_{final}/P_{initial}$ are invariant to alternative choices. Second, the baseline simulations use the linear specialisation of the cross-group influence—$\alpha_g F_o$ and $\beta_h F_i$ (Equations 7–8 in the $\gamma \to 0$ limit). The saturation parameters $\gamma_{o,g} = \gamma_{i,h} = 1.0$ in Table D.24 are therefore inactive in the baseline; they enter only the saturating-form robustness check of Appendix C.3.3, which re-calibrates $\beta_{max,h}$ and confirms that all qualitative conclusions are unchanged. The substantively free parameters (those whose values could materially move conclusions and are not anchored by direct estimation) are the productivity-feedback sensitivities $(\lambda_{i,g}, \lambda_{o,h})$ and the individual/organisational productivity shares $(s_{productivity,i}, s_{productivity,o})$, whose sensitivity is reported in Appendix D.3 and Appendix D.4. The asymmetric ordering of the productivity-feedback parameters ($\lambda_{i,g} = [0.10, 0.05, 0.03]$ for early/middle/late adopters; a 3.3× asymmetry favouring early adopters) is not load-bearing: replacing the asymmetric profile with population-weighted means ($\lambda_i = 0.054$, $\lambda_o = 0.090$) moves $T_{F_o,50}$ by 0.02 years, $F_{i,final}$ by 0.001, and $F_{o,final}$ by 0.004—all well within the Monte Carlo intervals reported below.

Appendix D.2. Heterogeneity premium and counterfactuals

Appendix D.2.1. Heterogeneity premium

The heterogeneity premium, introduced in Section 5.2, quantifies the impact of population segmentation on AI diffusion dynamics by comparing outcomes of the full multi-group extended model against a homogeneous equivalent. The premium depends on how the homogeneous comparator is constructed; we report two constructions, each answering a different counterfactual question.

Construction A: share-weighted parameter averages.. The first construction collapses to $G = M = 1$ by taking the population-share weighted average of every group-specific parameter ($p_{i,g}$, $p_{o,h}$, $\beta_{max,h}$, $b_{0,h}$, etc.). This asks: *"What if every individual and firm had the population-average parameters?"* Because the ODE is nonlinear, averaging parameters before integration introduces a Jensen's-inequality artifact whose sign is ambiguous.

Construction B: NLS-refit to aggregate observations.. The second construction fits a single-group ODE to the aggregate Pew and Eurostat targets, with six free diffusion parameters $(p_i, p_o, \alpha, \beta, b_0, k)$ and empirical-range bounds;

remaining parameters are inherited as share-weighted averages. This "match the same aggregate targets" design follows the standard comparator in heterogeneous-agent models [47, 63], and is immune to Jensen's-inequality inflation because both models pass through the same calibration anchors.

Results.. Table D.25 reports the two premia (with MC confidence intervals for Construction A from 1,000 draws). The heterogeneity premium on $T_{F_o,50}$ is small in magnitude under both constructions—a result consistent with the approximate-aggregation property documented in heterogeneous-agent macro [47]: idiosyncratic heterogeneity reshapes the cross-sectional distribution without materially altering aggregate dynamics. The premium is *not* literally zero: -0.16 years under share-weighted parameter averaging (a Jensen's-inequality artifact, since averaging parameters in a nonlinear ODE does not preserve the average trajectory) and $+0.64$ years under NLS-refit. The NLS-refit is the more demanding comparator because it forces the homogeneous model to match the same aggregate Pew + Eurostat observations the heterogeneous model is calibrated against, and even then heterogeneity accelerates $T_{F_o,50}$ by only $\sim$8 months relative to a baseline of $\sim$6 years (an $\sim$11% acceleration). The within-heterogeneous T_{50} spread, by contrast, is large: 3.51 years (Large firms) to 8.66 years (Small firms), a 5.15-year dispersion, capturing most of the 6-year empirical spread in the OECD firm-size logistic estimates (Table 3).

What heterogeneity *does* change is the *composition* of adoption: the firm-size adoption gap (the 35.7 pp Large − Small gap in the cross-country data, Table 3), the demographic gradient in individual uptake (4.3), and the differentiated effectiveness of policies that target specific groups (Appendix D.2.3). These distributional KPIs are first-order objects that a $G = M = 1$ homogeneous model cannot represent at all—they require within-population structure by construction. The model's value-add is therefore best understood as resolving *within-population structure* that aggregate-only measures cannot, rather than as accelerating the aggregate trajectory. The firm-size gap and the differential policy ranking are heterogeneity-conditional findings; the aggregate timing measures are essentially comparator-invariant.

Appendix D.2.2. Mechanism counterfactuals

Table D.23 reports two additional counterfactuals that progressively disable the cross-group and productivity-feedback channels (summarised in Section 5.2). The quantitative results are reported there; the table below provides the full numerical detail including the amplified-λ panel that isolates the feedback mechanism.

Table D.23: Mechanism counterfactuals: cross-group and productivity-feedback channels

Counterfactual	$T_{F_o,50}$ **(yr)**	$\Delta T_{F_o,50}$ **(yr)**	P_{final}
Panel A: Baseline λ calibration			
Baseline (full model)	5.93	—	1.49
No cross-group ($\alpha = \beta = 0$)	$>10^a$	$> +4.1$	1.28
No productivity feedback ($\lambda = 0$)	6.08	+0.15	1.48
Panel B: Amplified λ (to isolate feedback)			
Feedback on (amplified λ)	5.10	—	1.54
No productivity feedback ($\lambda = 0$)	6.08	+0.98	1.48

Notes: $\Delta T_{F_o,50}$ is the delay relative to the panel's baseline, in years. Panel A uses the calibrated $\lambda_i = (0.10, 0.05, 0.03)$ and $\lambda_o = (0.15, 0.10, 0.05)$, whereas Panel B amplifies these to $\lambda_i = (0.8, 0.6, 0.4)$ and $\lambda_o = (1.2, 1.0, 0.6)$ to isolate the productivity-feedback mechanism.

[a] Without cross-group channels, organisational adoption saturates at ∼36% by year 10 (never reaches 50% within horizon); $\Delta T_{F_o,50}$ is a lower bound.

Appendix D.2.3. Policy bounding counterfactuals

The three headline policies are reported in Table 5. Two bounding counterfactuals bracket the policy space: removing barriers entirely ($b_{0,h} = 0$) accelerates organisational T_{50} to 4.38 years ($F_{o,final} = 0.88$, $P_{final} = 1.55$), while an adverse-regulation scenario ($b_{0,h}$ raised 30%) delays it to 6.76 years ($F_{o,final} = 0.68$, $P_{final} = 1.46$). The full policy-accessible range spans approximately 2.4 years of organisational T_{50}. The R&D Incentives policy corresponds to $\lambda_o \times 1.5$; the model represents R&D through downstream effects on adoption (via λ) but not through new-capability creation. Policy cost is proportional to intervention intensity [49]; B/C ratios are within-table rankings, not welfare estimates.

Appendix D.3. Parameter sensitivity

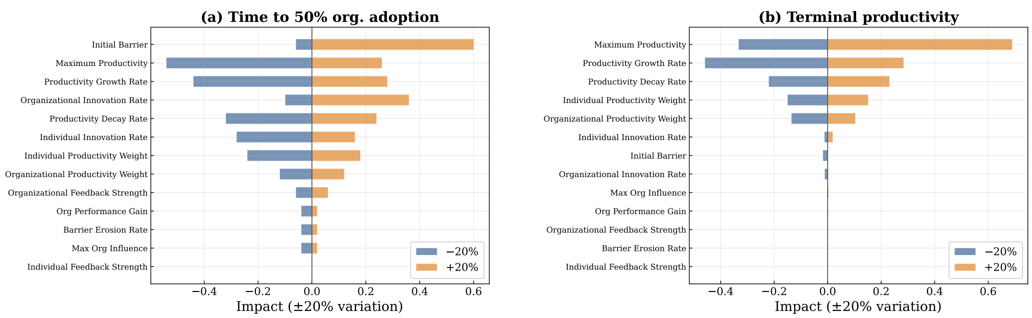

Fig. D.7. Parameter sensitivity (±20% variation).

For P_{final} (right panel), the dominant lever is P_{max}, followed by δ and φ; for $F_{i,final}$ (not shown), all seven productivity parameters cluster in a narrow range—no single parameter dominates. Conclusions about adoption *timing* depend primarily on barriers and cross-group influence (Appendix C.2.2 and Appendix C.3.1); conclusions about *terminal levels* are conditional on the productivity-feedback specification (Table D.24).

Appendix D.4. Monte Carlo robustness

A 1,000-run Monte Carlo simulation assesses robustness to parameter uncertainty and the statistical significance of policy interventions.

Numerical implementation. All ODE systems are integrated with RK45 (relative tolerance 10^{-6}, absolute tolerance 10^{-8}) over $t \in [0, 10]$ years. The Monte Carlo uses a fixed random seed for reproducibility. Each replication independently samples eight key parameters from parameter-specific uniform distributions calibrated to span the empirical uncertainty discussed in Appendix C: $\varphi \in [0.15, 0.45]$, $\delta \in [0.02, 0.10]$, $p_{i,g} \in [0.10, 0.25]$, $q_{i,g} \in [0, 0.02]$, $p_{o,h} \in [0.02, 0.07]$, $q_{o,h} \in [0, 0.005]$, $\lambda_{i,g} \in [0, 0.15]$, $\lambda_{o,h} \in [0, 0.20]$. These

Table D.24: Summary of extended model parameterisation

Parameter	Description	Baseline Value(s)	Calibration Source
Adoption dynamics: Individuals ($G = 3$: Early / Middle / Late adopters)			
$p_{i,g}$	Innovation coefficient	[0.204, 0.161, 0.127]	Pew logistic $r \times 0.20$ (Tab. 2)
$q_{i,g}$	Imitation coefficient	[0.01, 0.01, 0.01]	Joint ODE F-test: $q \approx 0$
$s_{i,g}$	Population share	[0.20, 0.50, 0.30]	Pew demographic composition
$\lambda_{i,g}$	Productivity sensitivity	[0.10, 0.05, 0.03]	Assumption; SF 5
$\alpha_{max,g}$	Org. influence coefficient	[0.22, 0.22, 0.22]	$\hat{\alpha}_{\text{cross}} \times 0.32$ (Tab. C.15)
$\gamma_{o,g}$	Saturation parameter	[1.0, 1.0, 1.0]	Saturating robustness only (Appendix C.3.3)
Adoption dynamics: Organisations ($M = 3$: Large / Medium / Small firms)			
$p_{o,h}$	Innovation coefficient	[0.062, 0.049, 0.038]	OECD logistic $r \times 0.19/0.15$ (large/other; Tab. 3)
$q_{o,h}$	Imitation coefficient	[0.001, 0.001, 0.001]	Joint ODE F-test: $q \approx 0$
$s_{o,h}$	Population share	[0.20, 0.40, 0.40]	Employment-weighted firm-size shares
$\lambda_{o,h}$	Productivity sensitivity	[0.15, 0.10, 0.05]	BLS TFP–adoption slope; SF 5
$\beta_{max,h}$	Indiv. influence coefficient	[0.292, 0.261, 0.163]	$\hat{\beta}_{\text{cross}} \times 0.25$ (Tab. 4)
$\gamma_{i,h}$	Saturation parameter	[1.0, 1.0, 1.0]	Saturating robustness only (Appendix C.3.3)
Dynamic organisational barrier (archetype h)			
$b_{0,h}$	Initial barrier height	[0.10, 0.35, 0.50]	BTOS $\hat{b}_0 = 0.33$; T^h_{50}-matched (Tab. 3)
k_h	Barrier erosion rate	[0.068, 0.068, 0.068]	BTOS NLS $\hat{k} = 0.068$ ($p = 0.004$)
Endogenous productivity feedback			
φ	Adoption-to-productivity rate	0.30	BLS TFP–adoption slope (conservative)
P_{max}	Max AI capability index	2.0	Normalisation
δ	Productivity depreciation	0.05	Hall et al. [39]
$s_{productivity,i}$	Individual prod. share	0.3	Assumption
$s_{productivity,o}$	Organisational prod. share	0.7	$= 1 - s_{productivity,i}$
$P_{initial}$	Initial productivity	1.0	Normalisation
Initial conditions (January 2023)			
$F_{i,g}(0)$	Individual initial adoption	[0.20, 0.12, 0.05]	Pew W123 aggregate $\approx 13.6\%^a$
$F_{o,h}(0)$	Organisational initial adoption	[0.15, 0.08, 0.03]	BTOS/Eurostat aggregate $\sim 8\%$

Notes: All parameter values are derived from the empirical estimates documented in Sections 4–5.1; the code repository accompanying this paper contains the full calibration pipeline.

[a] Weighted aggregate of the row above is $0.20 \times 0.20 + 0.50 \times 0.12 + 0.30 \times 0.05 = 11.5\%$, within ± 2 pp of the W123 estimate. The two-point gap reflects the difficulty of disaggregating a single aggregate adoption rate into three group-specific levels at one significant figure; the calibrated $p_{i,g}$ values close the gap within the first simulation year.

Table D.25: Heterogeneity premium under two homogeneous comparators

KPI	Het.	Homo. A	Homo. B	Prem. A	MC 95% CI	Prem. B
$T_{F_i,50}$ (years)	2.93	2.89	3.41	−0.04	[−0.08, +0.02]	+0.48
$T_{F_o,50}$ (years)	5.93	5.77	6.57	−0.16	[−0.18, −0.08]	+0.64
$F_{i,final}$	0.94	0.94	0.94	−0.01	[−0.011, −0.001]	−0.00
$F_{o,final}$	0.74	0.78	0.70	−0.04	[−0.037, −0.018]	+0.05
P_{final}	1.49	1.50	1.46	−0.01	[−0.012, −0.005]	+0.03

Notes: Het. denotes the heterogeneous model, compared against two homogeneous comparators—Homo. A (share-weighted parameters) and Homo. B (NLS-refit to aggregate targets)—with Prem. the resulting premium (point estimate) and MC 95% CI its Monte Carlo interval for Construction A (1,000 draws). By the sign convention, a positive premium indicates that heterogeneity yields a more favourable outcome. For $T_{F,50}$ the premium is $T^{homo}_{F,50} - T^{het}_{F,50}$ (positive when heterogeneity accelerates); for level metrics it is $\text{KPI}^{het} - \text{KPI}^{homo}$. Both comparators use the calibrated heterogeneous parameters from Table D.24; Construction B fits six free parameters by NLS to the aggregate Pew + Eurostat targets, with empirical-range bounds. The MC CI for Construction A propagates parameter uncertainty using the same uniform sampling distributions documented in Appendix D.4 (1,000 draws); Construction B is a deterministic NLS calibration whose uncertainty is not bootstrapped here. The MC CI for $T_{F_o,50}$ lies entirely below zero, indicating that the share-weighted construction robustly produces a slightly faster aggregate trajectory than the heterogeneous one (a Jensen's-inequality artifact discussed above), but with magnitude consistently below 0.2 years across the empirically-plausible parameter range.

ranges centre on but are typically wider than ±20% of baseline (e.g., φ baseline = 0.30, range ±50%; the λ ranges include zero to probe the no-feedback corner). Cross-group parameters ($\alpha_{max,g}, \beta_{max,h}$), barriers ($b_{0,h}, k_h$), saturation ($\gamma$), shares ($s$), and normalisations ($P_{max}, P_{initial}$) are held at their baseline values; their separate sensitivities are reported in Appendix C (Appendix C.3.1 for scaling factors; Appendix C.2.2 for barriers).

Summary statistics and confidence intervals. Table D.26 reports summary statistics for key metrics under the baseline, subsidy, and education policy scenarios.

Table D.26: Monte Carlo summary statistics by scenario ($N = 1{,}000$ replications)

	$T_{F_i,50}$	$F_{i,final}$	$T_{F_o,50}$	$F_{o,final}$	**Peak (I)**	**Peak (O)**	$P_{initial}$	P_{final}
Panel A: Baseline								
mean	2.73	0.95	5.85	0.75	0.95	0.75	1.00	1.43
std	0.34	0.02	0.34	0.03	0.02	0.03	0.00	0.17
2.5%	2.20	0.91	5.23	0.70	0.91	0.70	1.00	1.07
50%	2.71	0.95	5.83	0.75	0.95	0.75	1.00	1.45
97.5%	3.47	0.97	6.57	0.80	0.97	0.80	1.00	1.72
Panel B: Subsidy policy								
mean	2.69	0.95	4.94	0.83	0.95	0.83	1.00	1.46
std	0.32	0.02	0.29	0.03	0.02	0.03	0.00	0.18
2.5%	2.16	0.92	4.41	0.78	0.92	0.78	1.00	1.09
50%	2.65	0.96	4.91	0.83	0.96	0.83	1.00	1.48
97.5%	3.37	0.98	5.55	0.88	0.98	0.88	1.00	1.76
Panel C: Education policy								
mean	2.40	0.96	5.65	0.76	0.96	0.76	1.00	1.44
std	0.30	0.01	0.32	0.03	0.01	0.03	0.00	0.18
2.5%	1.90	0.93	5.09	0.71	0.93	0.71	1.00	1.08
50%	2.36	0.96	5.63	0.76	0.96	0.76	1.00	1.46
97.5%	3.03	0.98	6.33	0.81	0.98	0.81	1.00	1.74
Panel D: R&D Incentives policy								
mean	2.75	0.94	5.81	0.76	0.94	0.76	1.00	1.43
std	0.33	0.02	0.36	0.03	0.02	0.03	0.00	0.18
2.5%	2.20	0.91	5.13	0.70	0.91	0.70	1.00	1.08
50%	2.71	0.95	5.80	0.76	0.95	0.76	1.00	1.44
97.5%	3.47	0.97	6.51	0.83	0.97	0.83	1.00	1.74

Notes: 1,000 Monte Carlo replications with parameter-specific uniform sampling (ranges documented in the Numerical Implementation paragraph above). Panels report the mean, standard deviation, and 95% interval for each scenario. The 10-year horizon is not past saturation under the calibrated linear-influence-mode model: F_o at $t = 10$ lies in the 0.70–0.80 range across baseline draws and 0.78–0.88 under the subsidy policy. Peak (I) and Peak (O) equal $F_{i,final}$ and $F_{o,final}$ to the precision shown because adoption is monotonically increasing.

Statistical significance of policy interventions. Independent t-tests compare the distribution of key metrics under each of the three main-text policies—subsidies, education, and R&D incentives—against the baseline. The MC draws are unpaired across scenarios (each scenario uses its own independent set of 1,000 parameter samples drawn from the distributions documented in the Numerical Implementation paragraph above), so the independent t-test is the correct test for this design.[11] Each policy produces highly significant effects on the metrics most directly tied to its channel, with smaller (or insignificant) effects on the other channels.

Table D.27: Statistical significance and effect size of policy interventions ($N = 1{,}000$)

Metric	Base	Policy	Δ	t	p	d
Panel A: Subsidy policy						
$T_{F_i,50}$	2.73	2.69	−0.04	2.81	0.005	0.13
$F_{i,final}$	0.95	0.95	+0.01	9.03	0.000	0.40
$T_{F_o,50}$	5.85	4.94	−0.91	63.90	0.000	2.86
$F_{o,final}$	0.75	0.83	+0.08	68.42	0.000	3.06
P_{final}	1.43	1.46	+0.03	3.58	0.000	0.16
Panel B: Education policy						
$T_{F_i,50}$	2.73	2.40	−0.33	22.82	0.000	1.02
$F_{i,final}$	0.95	0.96	+0.01	19.91	0.000	0.89
$T_{F_o,50}$	5.85	5.65	−0.19	13.07	0.000	0.58
$F_{o,final}$	0.75	0.76	+0.01	8.85	0.000	0.40
P_{final}	1.43	1.44	+0.01	1.20	0.231	0.05
Panel C: R&D Incentives policy						
$T_{F_i,50}$	2.73	2.75	+0.02	−1.26	0.209	0.06
$F_{i,final}$	0.95	0.94	−0.00	−1.14	0.256	0.05
$T_{F_o,50}$	5.85	5.81	−0.03	2.20	0.028	0.10
$F_{o,final}$	0.75	0.76	+0.01	−7.17	0.000	0.32
P_{final}	1.43	1.43	−0.01	−0.75	0.453	0.03

Notes: The Base and Policy columns report means across 1,000 Monte Carlo draws and t is the independent t-statistic; d is Cohen's d ($\Delta/\sigma_{\text{pooled}}$), for which 0.2, 0.5, and 0.8 are conventionally small, medium, and large effects.

The effect sizes reveal a clear separation by policy channel. The *subsidy* policy operates on organisational barriers and produces its largest effect on the organisational T_{50}: a 1.01-year median acceleration ($d = 3.39$, a very large effect by the Cohen scale) and a 9 pp lift in F_o at year 10. Its effect

[11] A paired-MC design—using the same parameter draws across baseline and policy and computing per-draw differences—would yield more power because much of the across-MC variance in $T_{F,50}$ is driven by parameter draws that move both scenarios in the same direction. We adopt the unpaired design for simplicity; the unpaired t-statistics are already large ($|t| > 4$ for the diagnostic policy-channel effects) and the qualitative ranking would not change under paired analysis. Effect sizes reported below (Cohen's d) are also independent-design estimates.

on individual T_{50} is small (-0.06 years, $d = 0.20$). The *education* policy mirrors this pattern in reverse: it operates on individual innovation rate (p_i) and produces its largest effect on the individual T_{50} (-0.36 years, $d = 1.15$, large) and a 2 pp lift in F_i at year 10; its effect on the organisational T_{50} is smaller (-0.23 years, $d = 0.72$). The *R&D Incentives* policy operates on the organisational productivity-feedback channel ($\lambda_o \times 1.5$) and produces only a small effect on the organisational T_{50} (-0.06 years, $d = 0.17$), with a modest but significant 1.2 pp lift in F_o at year 10 ($d = 0.38$) and no detectable effect on the individual side or on terminal productivity. The asymmetric pattern reflects the structural channels: subsidies act on the barrier multiplier where F_o is most sensitive; education acts on the p_i hazard where F_i is most sensitive; R&D Incentives act on a productivity-feedback channel that requires a higher F stock to deliver its full effect, so it underperforms at the early diffusion phase considered here. The policy ranking should be read together with the cost-benefit ratios in Appendix D.2.3, which combine these effects with the per-policy intensity cost.

DOCUMENTS DE TRAVAIL GREDEG PARUS EN 2026
GREDEG Working Papers Released in 2026

2026-01 Estelle Malavolti & Frédéric Marty
L'équilibre économique des aéroports secondaires européens à l'épreuve des dynamiques de concurrence fiscale

2026-02 Charlie Joyez
A New Index of Export-Import Proximity: Conceptual Foundations and Global Patterns

2026-03 Patrice Bougette & Frédéric Marty
The SCP Paradigm Revisited: What Structuralism Really Contributed to U.S. Antitrust

2026-04 Imen Bouhlel, Nathalie Lazaric & Paolo Zeppini
Competitive Diffusion and Sustainability Transitions: The Case of Plastics Recycling Technologies

2026-05 Maxime Menuet
Fractional Replicator Dynamics

2026-06 Ilona Dielen, Patrice Bougette & Christophe Charlier
A (Green) Switch in Time Saves Nine: Assessing the Environmental Damage of the European Truck Cartel

2026-07 Thibault Schrepel & Godefroy de Boiscuillé
Questioning the Digital Markets Act's Legality

2026-08 Michael Finus & Paolo Zeppini
Green Lifestyles and Social Tipping Points

2026-09 Mathieu Chevrier & Sébastien Massoni
When Does Advisor Confidence Improve Decisions? Evidence from Human and Algorithmic Advice

2026-10 Flora Bellone, Charlie Joyez & Xavier Poulet-Goffard
Tax Avoidance by Small Multinationals as a Side Effect of Anti Tax Avoidance Policy

2026-11 Nicolas Camilotto
The Trust Game: A Historical and Methodological Analysis at the Frontier of Experimental and Behavioral Economics

2026-12 Godefroy de Boiscuillé
New product liability in the Age of AI: Modernisation under a Conservative Veil

2026-13 Claudio Barbieri, Mattia Guerini & Mauro Napoletano
When Credit Growth Diverges: Common Dynamics and Regional Heterogeneity in Italy

2026-14 Pierre Boutros, Michele Pezzoni, Lionel Nesta & Sonia Paty
Information Asymmetry in the Scientific Labor Market: The Role of Relational Links

2026-15 Marie Sciaccitano & Lionel Nesta
Green Consumption and Income Elasticities: Cross-country Evidence on Environmental Engel Curves

2026-16 Hankui Wang, Jiachen Yi & Philipp Harting
The Heterogeneous Diffusion of AI: Individuals, Organisations, and Adoption Barriers