WORKING PAPER IDB-WP-1881

Uneven Returns from Digitalization and AI

Economic Mechanisms and Policy Implications

Razvan Vlaicu

Inter-American Development Bank
Department of Research and Chief Economist

November 2026

Uneven Returns from Digitalization and AI

Economic Mechanisms and Policy Implications

Razvan Vlaicu[†]

[†] Research Department, Inter-American Development Bank; Email: vlaicu@iadb.org

Inter-American Development Bank
Department of Research and Chief Economist

November 2026

JEL Codes: O33, O38, D24, O54
Keywords: digitalization, artificial intelligence, general-purpose technologies, productivity, economic returns, Latin America and the Caribbean

To explore more IDB publications, visit
https://publications.iadb.org/en

Copyright © **2026** Inter-American Development Bank (IDB). This work is subject to a Creative Commons Attribution 4.0 International Public License CC BY 4.0 (https://creativecommons.org/licenses/by/4.0/legalcode). The terms and conditions indicated in the URL link must be met and the respective recognition must be granted to the IDB.

Any and all disputes arising under this license that cannot be settled amicably shall be resolved in accordance with the following procedure. Pursuant to a notice of mediation communicated by reasonable means by either you or the licensor to the other, the dispute shall be submitted to non-binding mediation conducted in accordance with the World Intellectual Property Organization (WIPO) Mediation Rules. Any dispute that cannot be settled amicably shall be submitted to arbitration pursuant to the United Nations Commission on International Trade Law (UNCITRAL) rules. The use of the IDB name for any purpose other than the respective recognition and the use of the IDB logo are not authorized by this license and require an additional license agreement.

Note that the URL link includes terms and conditions that are an integral part of this license.

The opinions expressed in the work are those of its authors and do not necessarily reflect the views of the IDB, its Board of Executive Directors, or the countries they represent.



Abstract

Digitalization and artificial intelligence (AI) are increasingly viewed as general-purpose technologies with the potential to reshape economic activity, raise productivity, and improve public services. Yet their returns are highly uneven across firms, workers, and countries. This paper develops an economic framework for understanding these differences, focusing on four properties of digital and AI technologies: strong complementarities among inputs, the need for workflow and organizational redesign, economies of scale, and economies of scope. These features help explain why technological adoption alone may generate limited gains, why productivity effects often emerge with a lag, and why benefits tend to concentrate in firms and locations with stronger capabilities and infrastructure. The paper applies this framework to Latin America and the Caribbean and discusses policy priorities for expanding digitalization, strengthening digital government, addressing information and coordination failures, and broadening access to the productivity gains associated with AI.

The paper has benefited from feedback provided by academic advisors, editors, contributors, and participants in the DIA 2026 workshops. The findings and interpretations in this paper are those of the author and do not necessarily reflect the views of the Inter-American Development Bank or the governments it represents.

1 Introduction

Digitalization and artificial intelligence (AI) have moved from being supporting technologies to becoming central elements of economic strategy. Across advanced and developing economies alike, governments now describe digital connectivity, data, and AI capabilities as engines of productivity, resilience, and global competitiveness. Firms frame AI not only as a tool for automation but also as a pathway to new products, better decisions, and expanded markets. The speed of recent technological progress, especially in machine learning and generative AI, has sharpened expectations that a new wave of growth is within reach.

Yet the real-world pattern of returns is strikingly uneven. Some economies and firms appear to be capturing outsized benefits, while others face slower adoption, weaker productivity effects, or even rising inequality without commensurate gains. This divergence is visible across the global economy and is especially salient in Latin America and the Caribbean (LAC), a region grappling with long-standing productivity stagnation, high informality, and persistent skills gaps. Recent evidence suggests that even as AI use diffuses rapidly among consumers, the ecosystem that converts AI into scaled production and broad-based productivity remains highly uneven across countries in the region. For instance, a recent assessment from the Economic Commission for Latin America and the Caribbean (ECLAC) notes that while LAC represents about 6.6% of global gross domestic product (GDP), it attracts only around 1.12% of global AI investment, a gap that constrains the ability to scale innovation and industrial applications (ECLAC, 2025).[1]

This disconnect between adoption enthusiasm and investment capacity resembles earlier moments in the history of general-purpose technologies (GPTs), when the promise of transformation was clear but its macroeconomic impact emerged unevenly and often slowly. The computer revolution, for example, generated a "productivity paradox" in the late 20th century: firms invested heavily in information technology (IT), but aggregate productivity acceleration appeared only later and was concentrated among leading firms and sectors (Brynjolfsson et al., 2021; OECD, 2021). Similar concerns now shape debates about AI's aggregate impact. Acemoglu (2025) argues that the near-term macroeconomic gains from AI may be substantial but also limited by the scope of tasks that can be profitably automated or augmented under current organizational and institutional conditions. These debates underscore that technological potential is not the same as realized economic value.

[1] Mandon (2025) finds that national AI preparedness does not map mechanically into income or economic complexity: several middle- and lower-income countries outperform predicted readiness levels, with the relative importance of digital infrastructure, human capital, and regulatory and ethical frameworks varying across country groups.

The puzzle becomes sharper when we consider how the diffusion of digital tools has expanded dramatically. Globally, about 5.5 billion people, roughly 68% of the world population, were online in 2024, up from 65% in 2023, while 2.6 billion remained offline (ITU, 2024). This rapid expansion might suggest a broad-based foundation for digital productivity. But the quality of connections, affordability of data and devices, and distribution of advanced skills remain uneven, shifting the locus of inequality from mere access to "meaningful connectivity" (World Bank, 2021; ITU, 2024). In LAC, regional mobile coverage has expanded, yet substantial "usage gaps" persist. Many people live in areas with mobile coverage but do not use mobile internet, often because of cost, skills, or perceived relevance barriers (GSMA, 2023).

At the firm level, digitalization outcomes also diverge. Large firms, multinational subsidiaries, and "digital native" platforms often leverage advanced analytics, cloud computing, and AI to reconfigure operations, optimize supply chains, and personalize services. Meanwhile, many small and medium-sized enterprises (SMEs) struggle to move beyond basic adoption of websites or social media toward integrated data systems and AI-driven process improvements (IDB, 2022; OECD and CAF, 2023). This gap matters because SMEs dominate employment in much of LAC, meaning that uneven AI diffusion can translate rapidly into uneven income and productivity outcomes.

The public sector adds another dimension to this puzzle. Governments in the region increasingly experiment with digital identity systems, e-procurement platforms, and AI-assisted targeting of social programs. These initiatives can yield substantial efficiency and transparency gains. But such improvements require stable data governance, cybersecurity capacity, interoperable systems across ministries, and sustained investment in human capital. These conditions are uneven across countries and often weaker at subnational levels (OECD and CAF, 2023; World Bank, 2021).

Labor market implications further raise the stakes. An assessment by the World Bank and the International Labour Organization (ILO) finds that approximately 26%–38% of jobs in LAC could be exposed to generative AI in some way; 8%–14% of jobs may be "productivity enhanced" through transformation, while 2%–5% could face high automation risk with current capabilities. At the same time, digital infrastructure gaps may prevent up to 17 million jobs from leveraging potential productivity benefits (Gmyrek et al., 2024; World Bank and ILO, 2024). Recent evidence based on harmonized household surveys from 14 Latin American countries further shows substantial heterogeneity in exposure and displacement risk across workers, with women, younger, more educated, and formal workers often more exposed, while measures that explicitly account for task complementarity yield flatter gradients across the income and education distributions (Ciaschi et al., 2025). The implication is not simply that

AI may alter job tasks but that the capacity to benefit from AI is unequally distributed across workers, firms, and countries.

These patterns raise two related challenges: how firms and countries can maximize the returns from digitalization and AI investment, recognizing that the extent of digitalization itself varies enormously across LAC, and how these gains can be distributed more equitably across workers and geographies. These two challenges are related but distinct, and together they define the core policy question addressed in this paper. To examine these challenges, the paper develops an economic framework that treats digital and AI technologies as economic production systems with distinct inputs, outputs, and dynamic properties. These features are then connected to historical lessons from past GPTs and identify four core economic properties that shape both productivity potential and inequality of outcomes: strong complementarities, the need for workflow redesign, economies of scale, and economies of scope. Together, these properties help explain several important phenomena: the productivity timeline and J-curve pattern observed with earlier GPTs, the collapse of returns when any one complementary input is missing, and the geographic concentration of digital and AI activity in a small number of global hubs. Finally, the paper draws policy implications tailored to developing countries in LAC, emphasizing strategies that can both maximize returns and broaden their distribution across society.

2 Digitalization and AI as Production Systems

A useful way to clarify the economics of digitalization and AI is to adopt a "production function" perspective. In this view, digital and AI technologies transform a bundle of tangible and intangible inputs into outputs that include both marketable goods and intermediate services that improve efficiency, quality, and decision-making. This framing helps explain why these technologies can generate productivity growth in principle yet yield uneven returns in practice.

2.1 Defining Digital and AI Technologies

Digitalization refers to the adoption and use of digital technologies to transform processes, organizational structures, products, and services. It includes foundational layers such as broadband networks, cloud computing, enterprise software, sensors, and digital platforms, and it spans applications from e-commerce and digital payments to advanced manufacturing

and digital government (OECD, 2021; IDB, 2022). Digitalization thus represents both a technological shift and an organizational transformation.

AI technologies are a subset of digital technologies characterized by systems capable of performing tasks that typically require human intelligence, such as perception, pattern recognition, language understanding, and decision support, often by learning from data. Contemporary AI is dominated by machine learning (ML), including deep learning, and recently by large-scale foundation models that can be adapted to multiple tasks, including generating text, code, images, and audio (Stanford Institute for Human-Centered Artificial Intelligence, 2025; OECD, 2024).

While AI rests on digital infrastructure, its economic logic differs in important ways. As Agrawal et al. (2018, 2022) argue, many AI applications are best understood as machines that lower the cost of prediction. Prediction here is not only forecasting future events or outcomes; it also includes classifying images, detecting anomalies, ranking alternatives, and generating probabilistic recommendations. By reducing prediction costs, AI increases the value of complementary inputs such as data, domain expertise, and decision rights embedded in organizational processes.

2.2 Outputs: Digital and AI Production

Within a production function framework, the outputs of digitalization and AI can be grouped into four broad categories. First, digital systems produce data services, including the collection, storage, cleaning, and exchange of data across internal and external networks. The value of these services is often indirect: high-quality data enhance decision-making, supply chain coordination, and customer experience. Data services also underpin digital markets, where data flows are essential for online transactions, fraud detection, and logistics optimization (World Bank, 2021; OECD, 2021).

Second, AI produces classification and prediction services. These are the core outputs of ML systems, ranging from credit scoring in fintech and demand forecasting in retail to predictive maintenance in utilities. In LAC, the rapid expansion of digital finance illustrates this channel: platforms such as Nubank and Mercado Pago rely on large-scale data analytics to manage risk and personalize services, supporting inclusion among previously underserved consumers while raising regulatory and competition questions (IDB, 2021; OECD, 2023).

Third, digital technologies produce digital content and interfaces, including websites, applications, digital media, and online marketplaces. These outputs are important for market

access and are often the first stage of digitalization for SMEs. Global estimates indicate the scale of the digital market: UN Trade and Development (UNCTAD) estimates global e-commerce sales at roughly USD 26.7 trillion in 2019, with the majority concentrated in business-to-business transactions (UNCTAD, 2021).

Fourth, generative AI expands the frontier by producing synthetic content—text, code, images, and designs—that can be used as intermediate inputs across many industries. Experimental evidence shows that such tools can raise productivity for certain tasks, with particularly strong effects for less experienced workers when AI is integrated into well-defined workflows (Noy and Zhang, 2023; Brynjolfsson et al., 2025). More recent field evidence reinforces this pattern. Randomized trials with nearly 4,900 software developers find sizable increases in completed tasks from access to an AI coding assistant, again with larger gains among less experienced developers (Cui et al., 2026), while a quasi-experiment among programmers finds substantial increases in code output after access to a generative AI coding tool (Gambacorta et al., 2026). These findings suggest that generative AI can reshape knowledge-intensive services, including those where LAC countries are building export niches, but also underscore that measured gains depend on the task, the user, and the organizational setting.

2.3 Inputs: Tangible and Intangible Complements

The inputs into digital and AI production span multiple layers. On the tangible side, the most visible inputs include (i) compute hardware such as servers, graphics processing units, edge devices, and end-user equipment; (ii) telecommunications infrastructure including fiber networks, mobile broadband, submarine cables, and data centers; (iii) energy systems, which provide reliable and increasingly clean power for data centers and high-intensity compute workloads; and (iv) sensors and Internet of Things devices in manufacturing, agriculture, and logistics. These physical inputs are unevenly available across LAC. Data center density, reliable electricity, and advanced cloud availability remain concentrated in a few larger economies, reflecting both scale dynamics and investment constraints (IDB, 2022).

On the intangible side, returns increasingly depend on (i) software and algorithms, including open-source frameworks, proprietary models, and enterprise applications; (ii) data assets, both personal and non-personal, and the governance frameworks that enable their responsible use; (iii) human capital, ranging from basic digital literacy to advanced ML engineering and domain-specific AI application skills; and (iv) organizational capital, including management practices, workflow design, and the ability to redesign products and services around digital capabilities.

For the purposes of the model in this paper, it is useful to distinguish between task-specific human capital and experience, which AI may either complement or substitute for, and implementation-oriented skills such as data engineering, systems integration, and change management, which are more naturally treated as complements to AI adoption.

The weight of intangible inputs is critical to understanding uneven returns. The economic literature on IT and productivity emphasizes that digital investments generate stronger payoffs when paired with complementary organizational changes and skills; without these, technology may simply add complexity or create "islands of automation" that do not translate into system-wide efficiency gains (Brynjolfsson et al., 2021; Pisu et al., 2021).

2.4 How AI Differs from Earlier Digitalization

AI differs from previous digital technologies along several dimensions, but three are especially relevant for economic and policy analysis. First, AI exhibits learning and self-improvement capabilities. Traditional software follows rules; ML systems improve performance as they are trained on larger or higher-quality datasets. This characteristic creates feedback loops in which early adopters who can generate or access large datasets may benefit disproportionately, reinforcing first-mover advantages.

Second, AI, especially foundation models, enables reusability across tasks. A model trained on general language or multimodal data can be fine-tuned for customer service, legal drafting, education tools, or administrative processes. This potential for cross-domain reuse is a seed of economies of scope that are more powerful than those associated with earlier enterprise software. Third, AI increasingly acts as a general cognitive layer for organizations. Rather than simply digitizing existing processes, AI can change how decisions are made, by whom, and with what degree of autonomy. This amplifies the importance of governance, accountability, and safety, particularly in high-stakes public sector applications such as social protection targeting, border management, or healthcare triage.

In sum, a production function lens highlights that digital and AI technologies are not plug-and-play. They are complex systems whose outputs depend on complementary tangible and intangible inputs. This interdependence foreshadows why returns are likely to be uneven and why policy must focus not only on adoption incentives but also on the ecosystem that makes high-return use feasible.

3 Economic Framework: Adoption and Productivity

Economists have long used the concept of GPTs to describe innovations, such as steam power, electricity, and computing, that generate broad, economy-wide transformations. GPTs typically diffuse across many sectors, enable complementary innovations, and require significant adaptation of organizational practices to realize their potential (Bresnahan and Trajtenberg, 1995; Helpman, 1998). Digitalization and AI increasingly fit this profile, but they also introduce new features that amplify both promise and risk.

This section highlights four economic properties of the new technologies: strong complementarity of inputs, the need for workflow and business redesign, economies of scale driven by network effects, and economies of scope that are particularly salient for AI. Together, these properties explain why digital and AI technologies can raise productivity in both private and public sectors while widening gaps between leaders and laggards.[2]

3.1 A Firm-Level Model

The mechanisms described above can be summarized in a parsimonious firm-level model. The purpose of the model is not to provide a complete equilibrium theory of AI adoption but to make explicit how bottlenecks, substitution between AI and some worker skills, scale, scope, organizational adjustment, and heterogeneity can jointly generate threshold effects and divergence in realized returns. The model distinguishes between inputs that can be adjusted relatively quickly and capabilities that are slow-moving at the horizon of an adoption decision.

Consider firm i. Let a_i denote access to digital or AI services, and let h_i^T denote task-specific human capital or experience. These two inputs need not be perfect complements: for some tasks, AI can substitute for experience, while for others it may be complementary. Define task capability s_i and effective digital capability q_i as

$$q_i = \min\left\{s_i, \bar{d}_i, \bar{h}_i^I, \bar{r}_i\right\} \equiv \min\{s_i, b_i\}, \qquad s_i = \left[\omega a_i^\rho + (1-\omega)(h_i^T)^\rho\right]^{1/\rho}, \tag{1}$$

where $0 < \omega, \rho < 1$. The elasticity of substitution between AI and task-specific skills is $1/(1-\rho)$: positive values of ρ imply relatively high substitutability, while negative values imply stronger complementarity. The variables $\bar{d}_i$, $\bar{h}_i^I$, and $\bar{r}_i$ are the firm's available data and digital infrastructure, implementation-oriented human capital (such as data engineering and

[2]These properties are closely related to the economics of prior GPTs, which emphasizes complementarities, organizational adjustment, and increasing returns as central to diffusion and productivity effects.

change management capacity), and organizational readiness, respectively. These are treated as inherited capacities in the short run, and b_i denotes the tightest of these constraints (i.e., the bottleneck). Equation (1) therefore allows AI to substitute for some task-specific skills while preserving a genuine weakest link structure in the capabilities needed to implement it. If, for example, $\bar{h}_i^I = b_i$, implementation skill is the binding bottleneck and has a positive shadow value, whereas relaxing another nonbinding capacity does not raise q_i on the intensive margin. This interpretation is close to the binding constraint logic emphasized by Hausmann et al. (2008).

At given input prices, the constant elasticity of substitution (CES) nest implies a minimum variable cost c_i of producing one unit of s_i. This unit-cost index can be written as $c_i = c(p_{a,i}, p_{h,i}^T; \omega, \rho)$. Allowing c_i to vary across firms introduces cost-side heterogeneity directly: differences in AI service prices, wages for task-specific skills, financing conditions, or access to specialized providers can alter both adoption intensity and the adoption threshold. Policies that expand broadband or implementation skills can instead relax b_i, while implementation assistance can reduce the fixed cost introduced below.

The same core capability can often be reused across several activities. To distinguish scope from simple firm size, let $j = 1, \ldots, J$ index a common set of potential applications; let $\nu_{ij} \geq 0$ be the share of firm activity associated with application j, with $\sum_j \nu_{ij} = 1$; let $z_{ij} \in [0, 1]$ measure the feasibility of using the common digital capability in that activity; and let $\theta_{ij} > 0$ measure its value per unit of activity. Scope opportunities can be summarized as

$$\Theta_i = 1 + \kappa \sum_{j=1}^{J} \nu_{ij} z_{ij} \theta_{ij}, \qquad \kappa > 0. \tag{2}$$

Because the activity weights are normalized, Θ_i does not rise mechanically with firm size; N_i below captures scale separately. Scope increases when a larger share of economically important activities can reuse the same data, software, or organizational architecture, or when those applications become more valuable. The feasibility terms z_{ij} should also not be interpreted as technologically fixed. Workflow redesign can increase them by making previously isolated activities compatible with a common digital architecture. This interpretation is consistent with the literature showing that firms reorganize internal hierarchies and layers in response to changes in technology and market opportunities (Garicano, 2000; Caliendo and Rossi-Hansberg, 2012; Caliendo et al., 2015; Brambilla et al., 2025).

Let N_i measure the scale of the organization or its relevant user base, and let $A_i > 0$ summarize features of the surrounding environment that affect the payoff to productive use.

More explicitly, one can write $A_i = \exp(\gamma' Z_i)$, where Z_i is a vector that includes institutional and market characteristics and the elements of γ can be positive or negative. Thus, the model does not assume that every institutional feature raises adoption. For example, data governance and contract enforcement may raise returns, while informality or weak appropriability can lower them; labor costs operate primarily through c_i, while labor market institutions can affect either costs or the payoff to adoption depending on the channel.[3]

Let β summarize the discounted value of future gains. The gross benefit from effective digital capability is

$$B_i(q_i) = \beta A_i N_i^{1+\eta} \Theta_i q_i^{\alpha}, \qquad 0 < \alpha < 1, \quad \eta \geq 0. \tag{3}$$

In equation (3), η captures scale and network effects, while Θ_i captures reuse across a normalized portfolio of applications. The formulation therefore separates scale, scope, and the surrounding environment conceptually even though all three raise the payoff to capability. Without additional variation, these multiplicative components are not separately identified from the benefit function alone; the distinction is economic rather than an econometric identification claim.

Adoption also entails organizational costs. Let F_i denote the fixed cost of experimentation, integration, governance, and redesign, and let $\chi_i q_i$ denote the monetized one-time disruption cost associated with implementation at intensity q_i. Defining $\tilde{c}_i = c_i + \chi_i$, the firm's adoption problem is

$$\max_{0 \leq q_i \leq b_i} \Pi_i(q_i) = \beta A_i N_i^{1+\eta} \Theta_i q_i^{\alpha} - \tilde{c}_i q_i - F_i \mathbf{1}\{q_i > 0\}. \tag{4}$$

Equation (4) puts the short-run J-curve cost inside the firm's decision rather than appending it after adoption. It also makes clear that policy can operate on several margins: lowering c_i, reducing F_i, or relaxing the bottleneck b_i need not have the same effect.

Ignoring the bottleneck momentarily, the interior optimal intensity is q_i^u. The actual choice is

$$q_i^* = \min\{q_i^u, b_i\}, \qquad q_i^u = \left[\frac{\alpha \beta A_i N_i^{1+\eta} \Theta_i}{\tilde{c}_i} \right]^{\frac{1}{1-\alpha}}. \tag{5}$$

Equation (5) makes the bottleneck distinction explicit. If $q_i^u < b_i$, a lower variable cost raises adoption intensity. If $q_i^u > b_i$, the capacity b_i binds: lowering the cost of a scalable input raises profits but does not increase q_i^* until the bottleneck itself is relaxed. This is the

[3] This is consistent with the regional evidence discussed by Ciaschi et al. (2025), who emphasize that the effects of AI in LAC depend partly on labor market characteristics such as wage levels, informality, and labor market institutions. These factors can affect both firms' incentives to adopt labor-saving technologies and workers' ability to benefit from AI complementarity, so their effects on adoption and economic returns need not operate in a uniform direction.

intensive margin sense in which investment in a nonbinding input can have little immediate effect even though it can still affect the extensive adoption decision.

The firm adopts only if the surplus at q_i^* covers the fixed cost. The adoption threshold is

$$F_i \leq F_i^* \equiv \begin{cases} (1-\alpha)\beta A_i N_i^{1+\eta} \Theta_i (q_i^u)^\alpha, & q_i^u \leq b_i, \\ \beta A_i N_i^{1+\eta} \Theta_i b_i^\alpha - \tilde{c}_i b_i, & q_i^u > b_i. \end{cases} \tag{6}$$

Equation (6) preserves the model's central threshold result but gives it a sharper bottleneck interpretation. A price subsidy can increase F_i^* even when a bottleneck prevents an intensive margin response, while infrastructure, implementation training, or organizational reform can raise b_i directly. Coordinated policies can therefore be more effective than isolated interventions when several constraints are close to binding.

Cross-firm heterogeneity can be introduced without changing the basic logic. Let $X_i = (A_i, N_i, \Theta_i, b_i, \tilde{c}_i)$, and suppose the fixed cost F_i has conditional distribution $G_F(\cdot \mid X_i)$. Then the adoption rate among firms with characteristics X_i is $G_F(F_i^*(X_i) \mid X_i)$. More generally, the joint distribution of A_i, N_i, Θ_i, b_i, $\tilde{c}_i$, and F_i maps the single-firm threshold into an economy-wide distribution of adoption and returns. Policies that shift different parts of this joint distribution can therefore have different efficiency and distributional effects.

When the bottleneck is not binding, the amplification created by scale and scope can be seen directly from equation (5). The relevant elasticities are

$$\frac{\partial \ln q_i^u}{\partial \ln N_i} = \frac{1+\eta}{1-\alpha}, \qquad \frac{\partial \ln q_i^u}{\partial \ln \Theta_i} = \frac{1}{1-\alpha}. \tag{7}$$

Because $0 < \alpha < 1$ and $\eta \geq 0$, both elasticities exceed one in the unconstrained region. Once b_i binds, however, the corresponding local intensive margin elasticities of q_i^* are zero until the bottleneck is relaxed. Thus, scale and scope can amplify differences among unconstrained adopters, while bottlenecks create flat regions in which additional demand or cheaper scalable inputs do not immediately increase effective capability.

To connect the adoption problem to the productivity J-curve, let $y_{i,t}$ denote productivity per unit of organizational scale and let $\Delta y_{i,t}$ denote the change relative to the pre-adoption level. Define the recurring productivity gain after implementation as

$$g_i(q_i) = A_i N_i^\eta \Theta_i q_i^\alpha. \tag{8}$$

The gross benefit in equation (3) can then be interpreted as the discounted value of this

productivity gain applied across the organization since $B_i(q_i) = \beta N_i g_i(q_i)$.[4]

The one-time adjustment cost $\chi_i q_i$ in equation (4) has a direct productivity interpretation. Expressed per unit of organizational scale, a reduced-form J-curve is

$$\Delta y_{i,t} = \begin{cases} -(\chi_i/N_i)q_i, & t = 0, \\ g_i(q_i), & t \geq 1, \end{cases} \qquad \chi_i > 0. \tag{9}$$

Equation (9) captures the productivity J-curve: measured productivity temporarily falls during implementation and subsequently rises once the organization has adapted and the digital capability becomes productive.[5] Because χ_i also enters $\tilde{c}_i$ in equation (4), firms internalize this initial disruption when deciding whether and how intensively to adopt.

3.2 Further Economic Effects

Three simple extensions connect the firm-level threshold more directly to cumulative advantage, industry concentration, and labor market outcomes. First, adoption can create cumulative advantage. Let $k_{i,t}$ summarize reusable data and organizational capital. A minimal law of motion is

$$k_{i,t+1} = (1 - \delta_k)k_{i,t} + \phi N_i q^*_{i,t} \mathbf{1}\{F_i \leq F^*_i\}, \qquad 0 < \delta_k < 1, \quad \phi > 0. \tag{10}$$

Higher $k_{i,t}$ can relax future bottlenecks, raise application feasibility $z_{ij,t}$ in equation (2), or reduce future implementation costs. Firms that cross the adoption threshold therefore accumulate assets that make subsequent adoption easier, whereas persistent non-adopters may fall behind as those assets depreciate. This creates history dependence and gives temporary interventions potentially persistent effects when they move firms close enough to the threshold to begin accumulating k.

Cumulative advantage can also operate across firms. Let $m_{\ell,t}$ be the share of adopters in

[4]Formally, if the productivity gain is constant from period $t = 1$ onward and future gains are discounted at factor $0 < \delta < 1$, then $\beta = \sum_{t=1}^{\infty} \delta^t = \delta/(1 - \delta)$. Thus, β summarizes the present value of a perpetual stream of future gains.

[5]The term "productivity J-curve" is from Brynjolfsson et al. (2021), who use it to describe the pattern whereby investment in a new GPT is initially accompanied by weak measured productivity because firms are simultaneously accumulating complementary intangible assets, with productivity gains appearing later as those investments begin to pay off.

location or sector ℓ. A simple local externality is

$$A_{\ell,t} = \bar{A}_\ell \left(1 + \zeta m_{\ell,t}\right)^\lambda, \qquad \zeta \geq 0, \quad \lambda \geq 0. \tag{11}$$

Adoption can raise the local payoff to further adoption through specialized service providers, thicker skill pools, shared infrastructure, and knowledge spillovers. Equation (11) provides a formal basis for coordination failures and spatial clustering: when $\lambda > 0$, individually rational non-adoption can be self-reinforcing if firms expect few complementary adopters around them.

Second, the baseline benefit function is partial equilibrium. A standard CES demand extension makes the demand side explicit. Suppose a firm faces demand x_i, and digital capability lowers marginal production cost according to $m_i(q_i)$, where

$$x_i = D_i p_i^{-\varepsilon}, \qquad m_i(q_i) = m_{0i} \exp(-\tau q_i), \qquad \varepsilon > 1, \ \tau > 0, \tag{12}$$

and under monopolistic competition, the optimal price is

$$p_i = \frac{\varepsilon}{\varepsilon - 1} m_i(q_i). \tag{13}$$

This familiar CES structure (Dixit and Stiglitz, 1977) shows that adoption can raise a firm's market share while lowering its price. When adoption becomes widespread, some productivity gains are passed on to consumers through a lower sectoral price index, so private and social returns need not coincide. Combined with scale or data advantages, the same mechanism can reallocate market share toward early adopters and contribute to concentration even when the underlying markup formula is simple.

Third, a minimal labor market extension separates technological exposure from realized impact. For a worker in occupation o employed by firm i, let $e_o \in [0, 1]$ denote occupational exposure to AI and $c_o \in [0, 1]$ the degree to which exposed tasks are complementary with AI. Let $\psi(q)$ denote displacement pressure, with $g_i(0) = \psi(0) = 0$. A reduced-form wage or productivity effect can be written as

$$\Delta w_{oi} = e_o \left[c_o g_i(q_i^*) - (1 - c_o)\psi(q_i^*)\right]. \tag{14}$$

Equation (14) is in the spirit of complementarity-adjusted exposure measures proposed by Pizzinelli et al. (2023). It formalizes three points used later in the paper: exposure is not the same as impact, the effect of exposure depends on the capability of the worker's

employer, and weak digital infrastructure can buffer displacement in low-complementarity occupations while blocking augmentation in high-complementarity occupations. Taken together, equations (1)–(14) provide a compact framework for the mechanisms discussed next and the policy interventions considered in Section 5.

3.3 Economic Implications

The formal theoretical framework presented above captures several economic implications that arise from incorporating digital and AI technologies in production processes.

Strong Input Complementarity. Digital and AI technologies display strong complementarities among several implementation inputs, but this need not imply perfect complementarity between AI and every type of human capital. Equation (1) makes this distinction explicit. AI services and task-specific human capital h_i^T enter a CES nest and may be substitutes or complements depending on the task. By contrast, effective use of that task capability is constrained by slower-moving data and digital infrastructure, implementation-oriented human capital $\bar{h}_i^I$, and organizational readiness $\bar{r}_i$. These latter inputs can therefore operate as genuine bottlenecks. Consider tax enforcement. AI-assisted audits require digitized taxpayer records, consistent reporting standards, analytical and implementation capacity, and legal authority to act on algorithmic signals. If one of these elements is binding, expanding a different input need not raise effective capability until the bottleneck is relaxed.

This distinction also reconciles two strands of evidence that would otherwise appear inconsistent. Studies of digital productivity repeatedly show that returns rise when firms bundle technology adoption with management upgrades, workforce training, and process changes (Brynjolfsson et al., 2021; OECD, 2021). Recent firm-level evidence likewise finds stronger AI-related productivity gains when adoption is accompanied by software, data, training, or capable management (Aldasoro et al., 2026; Kazakis, 2026). At the same time, experimental evidence often finds larger gains for less experienced workers on particular tasks (Brynjolfsson et al., 2025; Cui et al., 2026). That pattern is consistent with substitution between AI and task-specific experience inside the CES nest, while data engineering, change management, data quality, and organizational readiness remain complementary implementation inputs. Experimental evidence also shows that gains depend strongly on whether tasks lie within the technology's capabilities and whether organizations provide appropriate guidance and workflow conditions (Dell'Acqua et al., 2026b).

These complementarities extend beyond the firm. A national AI strategy may fail without

competitive telecom markets, reliable energy, appropriate skills, usable data, and governance frameworks that permit responsible deployment. Moreover, some of these inputs are supplied by different actors. The local availability of technical service providers, training institutions, data standards, and digital public infrastructure can therefore determine the return to a firm's own investment. This interdependence is the basis for the coordination mechanism formalized in equation (11) and provides a rationale for public action focused on binding constraints rather than uniform subsidies to all digital inputs.

Workflow and Business Redesign. A second property is the need for deep workflow redesign. In earlier eras of IT, firms often digitized existing processes rather than rethinking processes around digital capabilities. The contrast between shallow digitalization and genuine workflow redesign is illustrated by permit issuance. A common approach simply moves existing paper forms online, leaving citizens as intermediaries between fragmented government systems. Larger gains arise when the process is redesigned: once a citizen is authenticated through a digital identity system, the permitting authority automatically retrieves information already held by the government. This eliminates redundant verification, enables parallel processing, and shifts the burden of coordination from citizens to government systems. AI intensifies the imperative for redesign. Because AI lowers the cost of prediction and classification, it changes the economics of decision-making. It can shift which tasks are best done by humans versus machines, and it can change how value is created in service industries.[6]

Recent field evidence is consistent with this distinction: access to generative AI can save workers time without automatically changing the broader composition of their tasks, suggesting that tool provision alone need not induce organizational transformation (Dillon et al., 2026). Thus, business transformation, not mere technology procurement, will determine long-run returns. Organizational structure is itself endogenous: firms add and remove managerial layers in response to technology and market shocks (Garicano, 2000; Caliendo and Rossi-Hansberg, 2012; Caliendo et al., 2015), and recent evidence from Chile shows firms reorganizing hierarchies in response to export opportunities (Brambilla et al., 2025). In the model, these adjustment and integration requirements are summarized by the fixed cost F_i, the disruption cost χ_i, and the application feasibility terms z_{ij} in equation (2).

An important qualification applies to much of the experimental evidence on generative AI. Many studies randomize access to a tool among workers in firms that have already agreed to experiment with it. These estimates are therefore especially informative about the intensive margin returns to tool access among selected organizations that are plausibly already above,

[6] As Agrawal et al. (2022) note, the main strategic challenge may not be obtaining AI but reorganizing the firm to use prediction effectively.

or close to, the adoption threshold; they are less informative about take-up among firms in the lower tail of the capability distribution.

The public sector faces analogous challenges. AI can help detect fraud, optimize inspection schedules, and personalize citizen services but only if agencies redesign data flows, ensure interoperability, and establish accountability frameworks. The Organisation for Economic Co-operation and Development (OECD) and the Development Bank of Latin America and the Caribbean (CAF) emphasize in their review of digital government in LAC that institutional coordination and governance structures are central to success and that fragmented systems often block cross-ministry benefits (OECD and CAF, 2023).

Economies of Scale and Network Effects. A third property is the presence of powerful economies of scale. Digital goods and AI models often have high fixed costs and near-zero marginal costs of replication. Platforms can spread infrastructure costs across large user bases. Data-driven learning can improve algorithms at low additional cost, while larger customer and user bases can raise the value of platforms through network effects. These forces can drive rapid market concentration and generate superstar firms. Consistent with this mechanism, recent firm-level evidence finds that measured productivity gains from AI adoption are concentrated among medium and large firms (Aldasoro et al., 2026). In the model, these scale advantages operate through N_i and the parameter η in equation (3). The CES-demand extension in equations (12)–(13) adds a market share channel: a productive adopter can lower its price and expand sales, while sector-wide adoption passes part of the gain to consumers. When scale advantages interact with local adoption externalities in equation (11), market shares and innovative activity can become increasingly concentrated even when firms begin with similar access to frontier technology.

LAC provides visible examples of these scale economies. Mercado Libre built regional logistics, payments, and advertising ecosystems that benefit from cross-market network effects, while fintech leaders have scaled digital credit and savings products using data-driven risk assessment. These firms can invest in AI more aggressively than smaller competitors, reinforcing an uneven competitive landscape. The same scale logic applies to cloud service markets and AI compute, where global providers benefit from massive capital pools and cross-border data pipelines.

Scale economies in digital and AI industries also have a spatial dimension. Because returns depend heavily on access to skilled labor, venture capital, research institutions, and dense networks of suppliers and collaborators, firms and talent tend to concentrate geographically. This agglomeration logic helps explain why a disproportionate share of global AI research,

investment, and firm formation has clustered in a small number of locations such as the San Francisco Bay Area and Seattle. For LAC, this dynamic is consequential. At the country level, economies of scale may favor large markets with strong internal demand, deeper capital markets, and robust research ecosystems. Smaller economies, including many Caribbean states, may need regional strategies to pool demand and share infrastructure to overcome scale disadvantages. Without deliberate policy to build regional innovation ecosystems, scale forces may continue to pull the most productive firms and workers toward established global hubs, compounding the region's existing investment gap.

Economies of Scope. A fourth property is economies of scope, which are particularly prominent for AI. Foundation models and shared data infrastructures allow the same core assets to be reused across multiple products and services. A government that builds a high-quality digital identity system, interoperable registries, and secure data-sharing frameworks can leverage these across taxation, social protection, health services, and disaster response. Similarly, a firm that builds a strong data architecture can apply AI to marketing, operations, compliance, and HR analytics.

These scope economies stem from several distinct sources: shared digital infrastructure such as cloud platforms, digital identity systems, and cybersecurity frameworks; reusable software and algorithms that can be adapted across applications; and interoperable data architectures that make information collected for one purpose available for others. Once these core assets are in place, deploying additional digital or AI-enabled services becomes relatively inexpensive. Recent firm-level evidence is consistent with this mechanism: productivity gains are stronger when AI is deployed across multiple business functions rather than in isolated applications (Pastor-Merino et al., 2026). This helps explain why governments with integrated data and software stacks can expand digital services across tax, social protection, health, and education at declining marginal cost and why firms with strong data architectures apply AI across marketing, operations, and compliance tasks. Equation (2) summarizes this channel through Θ_i. Because application weights are normalized by their share in firm activity, Θ_i is conceptually distinct from the scale variable N_i: it rises when common digital assets become usable in a broader or more valuable set of functions, not merely because the firm is larger. Moreover, the feasibility terms z_{ij} can change as the organization redesigns workflows, and the accumulation process in equation (10) can raise future scope by building reusable data and organizational capital.

Economies of scope also operate at the individual level. Once the fixed costs of digital adoption, including internet access, devices, and basic digital skills, are incurred for one use, the marginal cost of adopting additional digital services is very low. Widespread use of tools

such as WhatsApp lowers the incremental effort required to adopt other digital applications, including government services, education platforms, or health tools. The OECD's TALIS 2024 survey provides a suggestive illustration: reported AI use among lower secondary teachers varies widely across countries, and adoption in some LAC systems, including Colombia, exceeds that reported in the United States (OECD, 2025). Although such comparisons are descriptive and should not be interpreted causally, they show that AI uptake need not map mechanically into national income levels and may build on pre-existing patterns of digital use.

Productivity Timeline and the AI "J-curve." These four properties imply significant productivity potential but also suggest that benefits may take time to materialize. Brynjolfsson et al. (2021) argue that digital technologies often produce a J-curve pattern, in which productivity may initially stagnate or even decline as organizations incur adjustment costs, invest in intangible complements, and experiment with new workflows. Only later does productivity accelerate when complementary innovations diffuse and organizations learn how to scale effective practices. The still-limited evidence from the generative AI era is consistent with separating task-level gains from organization-wide effects: large experiments document substantial gains on particular tasks, while broader workplace studies often find smaller immediate changes in task composition or aggregate labor outcomes (Cui et al., 2026; Dillon et al., 2026; Humlum and Vestergaard, 2025a). Equation (9) captures this timing distinction in reduced form, while equation (4) makes the initial disruption cost part of the firm's adoption decision. The accumulation mechanism in equation (10) adds persistence: implementation today can build data and organizational capital that raise future returns, so the path after adoption need not simply return to a fixed steady state.

AI is likely to follow a similar path, with two important twists. The first is that AI's learning capabilities may accelerate performance improvements once systems are in place. The second is that regulatory uncertainty and trust concerns may slow deployment in high-stakes domains. Thus, plausible scenarios for AI productivity in LAC range from rapid micro-level gains in specific sectors (e.g., customer service, digital finance, and targeted logistics) to longer lags in broader public sector transformation, where procurement rules, data fragmentation, and skills constraints are more binding.

These dynamics imply that the relationship between investment and returns is likely to be non-linear. Below certain thresholds of connectivity, skills, or organizational readiness, additional technology investment may yield negligible returns. Once those thresholds are crossed, however, returns can accelerate rapidly as complementary inputs reinforce each other and network effects compound. This non-linearity has important implications for LAC.

Patience and ecosystem-building matter. Countries and firms that treat AI as a short-term cost-cutting mechanism may underinvest in the complementary capabilities required for exponential growth. Conversely, those that invest in data quality, skills, organizational redesign, and governance may move earlier to the steep upward part of the J-curve.

4 Why Economic Returns Diverge

We can now return to the motivating question. The four economic properties outlined above do not operate in isolation. In practice, their interactions generate a set of distinct mechanisms that help explain the uneven returns documented in Section 1. Strong complementarities create thresholds below which access to technology produces little economic value. The need for workflow redesign raises the fixed cost of crossing those thresholds and delays the realization of productivity gains. Economies of scale reward firms, countries, and platforms that can spread fixed digital investments over larger markets, while economies of scope reward organizations that can reuse data, software, models, and organizational capabilities across many applications. These mechanisms reinforce one another: firms with stronger initial capabilities can adopt earlier, learn more quickly, apply the technology more broadly, and spread the resulting fixed costs over larger volumes of activity.

The resulting divergence is therefore not simply a difference in the speed at which otherwise identical firms adopt the same technology. It reflects cumulative differences in complementary capabilities, organizational readiness, scale, and opportunities for reuse. These differences can operate across firms, sectors, countries, and territories, and they can also shape which workers are positioned to benefit from technological change. The subsections below examine these channels in turn and connect the framework developed in Section 3 to the empirical literature.

4.1 Complementarity as a "Capability Filter"

Because digital and AI technologies require multiple complementary inputs, adoption becomes a capability filter. High-return use requires not only access to the technology itself but also reliable connectivity, appropriate hardware and software, usable data, skilled workers, managerial capacity, and organizational changes. Where one of these inputs is seriously deficient, increasing the others may generate little additional output. This is the mechanism represented by equation (1). Equation (5) distinguishes an unconstrained firm from one

whose data, implementation skills, or organizational readiness is binding. Equation (6) then shows that a reduction in variable cost can still raise the probability of adoption even when it does not relax the binding capacity constraint, whereas infrastructure or implementation policies can raise the bottleneck b_i directly. The idea is closely related to models of technology adoption in which complementary technologies must be accumulated in the presence of fixed adjustment costs (Jovanovic and Stolyarov, 2000). In such settings, complementarities can produce discrete adoption decisions and substantial differences in observed technology use even when firms face similar frontier technologies.

The empirical literature on earlier waves of IT provides substantial support for this interpretation. Digital capital has generally produced larger productivity gains when combined with organizational capital, worker skills, and management improvements rather than deployed as a stand-alone input (Brynjolfsson et al., 2021; Pisu et al., 2021). Recent evidence on AI points in the same direction. Using a purpose-designed Census Bureau survey of U.S. manufacturing establishments, McElheran et al. (2026) find that industrial AI adoption is strongly associated with recent digital infrastructure, structured production process management, firm size, and expertise. Importantly, widespread digitization by itself does not imply intensive AI adoption. European firm-level evidence similarly indicates that the productivity effects associated with AI are larger where firms possess complementary software, data, training, and organizational capabilities (Aldasoro et al., 2026; Kazakis, 2026). Taken together, these findings support the interpretation of AI readiness as a bundle of capabilities rather than a single technology access variable.

This mechanism is particularly relevant for LAC because the distribution of these complements is highly uneven. Internet coverage and basic digital infrastructure have improved, but more advanced capabilities remain concentrated. Differences in higher-education quality, technical skills, R&D capacity, reliable infrastructure, access to finance, and institutional quality create very different environments for productive adoption. The Oxford Insights Government AI Readiness Index, for example, documents substantial differences within the region, with Brazil, Chile, and Uruguay among the strongest regional performers (Oxford Insights, 2024). More broadly, the distinction between basic connectivity and the infrastructure, data institutions, and human capital needed to use digital technologies productively is consistent with the broader development literature on digital dividends (World Bank, 2021).

The same filtering process operates within countries. Large and technologically sophisticated firms are generally better positioned to combine cloud infrastructure, proprietary data, specialized workers, and formal management systems. SMEs often face several constraints simultaneously. The Inter-American Development Bank's (IDB) evidence on firm digitalization

in LAC emphasizes that advanced adoption tends to build on pre-existing digital maturity (IDB, 2022). A small firm that lacks reliable records, skilled managers, standardized processes, and access to technical expertise may receive little return from acquiring an advanced AI application in isolation. Moreover, managers may not know which technology is appropriate for the firm's production process or how to reorganize complementary tasks around it. Experimentation itself is costly, while part of the information generated by successful experimentation can spill over to competitors.

Complementarity therefore creates a mechanism through which initially modest differences in capabilities can become persistent differences in digital performance. Organizations already near the threshold in equation (6) can respond strongly to declining AI costs, whereas those far below it may exhibit little response. The diffusion of inexpensive generative AI tools can consequently coexist with large gaps in economically meaningful adoption. In this sense, the relevant digital divide increasingly concerns not simply access to a technology but the ability to assemble the bundle of complementary assets required to turn that technology into sustained productivity gains. Equation (10) adds a dynamic implication: once adoption begins, use itself can build reusable data and organizational capital, making the gap between adopters and persistent non-adopters cumulative rather than purely static.

4.2 Workflow Redesign As Handicap for Late Adopters

The second mechanism is organizational. Digital and AI technologies often create their largest gains not by performing an existing task somewhat more cheaply but by changing how tasks are combined, sequenced, assigned, and monitored. Firms that simply place a new technology on top of an unchanged production process may therefore capture only a small share of its potential value. In the model, these experimentation, integration, and reorganization requirements appear in the fixed cost F_i, the implementation disruption χ_i, and the application feasibility terms z_{ij}. Because redesign can both raise current implementation costs and expand the future set of profitable uses, organizational readiness can generate large differences in the timing, scope, and magnitude of returns.

This distinction is consistent with the emerging experimental evidence on generative AI. Dell'Acqua et al. (2026b) show that AI can substantially improve knowledge-worker performance on tasks that lie within the technology's capabilities but can also reduce performance when users apply it outside those capabilities. Their result highlights the importance of task selection, managerial guidance, and the allocation of decision rights rather than access to the tool alone. Likewise, Dillon et al. (2026), studying thousands of workers across multiple

firms, find that individually assigned generative AI tools save time on activities such as email without automatically producing large changes in the broader composition of work. Individual productivity improvements therefore need not translate immediately into organization-wide transformation.

Recent experimental evidence on teamwork strengthens this point. Dell'Acqua et al. (2026a) study professionals working on product development problems and show that generative AI can alter how expertise is combined within organizations: individuals assisted by AI can perform at levels comparable to conventional teams on some dimensions, while AI also changes the integration of specialized knowledge and the organization of collaborative work. Such results suggest that the productivity consequences of AI extend beyond substitution between a worker and a machine. They concern the architecture of teams, the division of labor, and the role of human judgment. Those gains require organizations to reconsider workflows rather than merely distribute AI accounts to existing employees.

The organizational barrier is likely to be more binding where production processes are informal or weakly standardized. In many traditional sectors in developing countries—including retail, construction, smaller-scale manufacturing, tourism, and parts of agriculture—data may not be systematically collected, responsibilities may not be formally codified, and IT systems may consist of disconnected applications. In such settings, it is difficult to automate or augment a process that has not first been made explicit. Firms with stronger management practices can identify suitable use cases, establish performance metrics, experiment with alternative workflows, and learn from unsuccessful implementation. Firms with weaker management capacity may instead add digital tools without removing redundant processes, generating additional complexity rather than productivity.

The same logic applies to the public sector. Moving a paper form online is fundamentally different from redesigning a public service around interoperable data and automated information exchange. Digital identity, shared registries, and secure data access can allow an agency to eliminate repeated verification and reorganize a service around information already held by government. The literature on digital government in LAC stresses precisely these coordination and interoperability problems (OECD and CAF, 2023; Cristia and Vlaicu, 2022). Agencies with strong digital leadership and common data architectures can redesign tax administration, licensing, procurement, or benefit delivery, while those that retain fragmented databases and rigid administrative procedures may realize only superficial gains from digitization.

Workflow redesign thus provides a second source of cumulative advantage. Organizations that are already well managed face lower effective adjustment costs, learn faster from

experimentation, and can reorganize tasks as technological capabilities change. Late adopters confront not only the cost of the new technology but also a backlog of organizational modernization. This mechanism helps explain why task-level evidence of substantial AI productivity gains can coexist with slower changes in aggregate firm productivity: the diffusion of a tool can be rapid, while the complementary redesign of organizations remains gradual.

4.3 Scale Economies and Platform Dominance

A third mechanism arises from scale. Many digital products require substantial fixed expenditures on software, data architecture, computing infrastructure, cybersecurity, model development, and organizational integration, while the marginal cost of extending the resulting service to another transaction or user can be comparatively low. Equation (3) captures this mechanism through organizational scale N_i and the parameter η, while equation (7) shows how scale can amplify differences in optimal digital intensity. This cost structure means that identical technological capabilities can generate very different private returns depending on the size of the market over which they can be deployed.

Scale advantages are reinforced in platform markets by network effects and data accumulation. The digital economics literature emphasizes that many platforms operate in multisided markets in which participation on one side raises the value of participation on another and in which larger activity levels can generate additional data for improving matching, prediction, personalization, or risk assessment (Goldfarb and Tucker, 2020). Data can therefore reinforce conventional fixed-cost economies: a larger platform can spread technology expenditures across more users while producing information that improves its services. Because data are partly nonrival, their productive use can itself generate increasing returns (Jones and Tonetti, 2020). These forces do not imply that all digital markets become monopolies, but they do create economic conditions under which advantages associated with early scale can become persistent.

Firm-level evidence on AI is consistent with this mechanism. Recent European evidence finds that AI adoption and measured productivity gains are concentrated disproportionately among larger firms (Aldasoro et al., 2026). Larger firms can support specialized technical teams, purchase enterprise-scale computing services, amortize implementation costs across many employees or customers, and experiment simultaneously with multiple applications. Small firms can increasingly purchase frontier AI services through the cloud, which reduces some technological fixed costs, but this does not eliminate fixed costs associated with data preparation, integration, governance, management attention, and process redesign. Scale

therefore continues to matter even as the direct price of access to foundation models falls.

These forces also operate across countries. LAC's share of global AI investment remains substantially below its share of global GDP (ECLAC, 2025). Large economies offer deeper pools of skilled labor, larger domestic demand, denser supplier and research networks, and greater opportunities to spread infrastructure costs. Smaller economies face a different calculation. A cloud region, advanced data center, specialized AI research facility, or domestic platform may not be economically viable when demand is fragmented across a small national market. This is especially relevant for Caribbean economies, where national scale may be insufficient to support some high-fixed-cost components of the digital ecosystem.

The spatial dimension reinforces these differences. Technology-intensive industries tend to cluster where specialized workers, venture capital, universities, suppliers, and sophisticated customers are already concentrated. Once a location reaches sufficient scale, these agglomeration forces can attract further investment and talent, while peripheral locations may lose both. In LAC, this can produce concentration not only between countries but also within them, as capital cities and large metropolitan areas capture a disproportionate share of high-productivity digital activity.

Scale therefore has an ambiguous economic implication. It can produce genuine efficiency gains: spreading fixed costs over a larger base lowers average cost and can finance better infrastructure and services. The demand extension in equations (12)–(13) also clarifies that adoption gains need not accrue entirely to firms: competitive responses and lower prices can transfer part of the surplus to consumers. At the same time, scale can become self-reinforcing when combined with network effects, proprietary data, switching costs, and ecosystem integration. The relevant explanation for uneven returns is therefore not simply that large firms possess market power. It is that the underlying economics of digital and AI production often make large user bases, transaction volumes, and data pools intrinsically valuable. These advantages can then interact with complementarities and organizational capabilities to widen the distance between leading firms and smaller or later-adopting competitors.

4.4 Scope Economies and the Advantages of Data-Rich Ecosystems

Economies of scope provide a fourth and distinct mechanism of divergence. Once an organization has built a high-quality data architecture, cloud infrastructure, cybersecurity systems, AI governance procedures, and a workforce able to use these assets, the same underlying capability can often be applied to many different tasks. Equation (2) captures this mechanism

through Θ_i, the value of the set of applications available to the organization. The greater the number and value of feasible uses, the larger the return to the common digital capability. Equation (7) consequently implies that organizations with broader opportunities for reuse will choose higher levels of effective digital capability.

The economics of data provides a particularly useful foundation for this argument. Unlike most physical inputs, data are substantially nonrival: using a dataset for one application does not necessarily prevent the same information from being used for another. Jones and Tonetti (2020) show that this characteristic can generate increasing returns and potentially large gains from broader productive use of data. The implication for AI is direct. A well-governed customer, transaction, logistics, or administrative database can support multiple prediction problems rather than a single application. Similarly, investments in data cleaning, identifiers, access protocols, and interoperability can lower the cost of subsequent applications because much of the enabling infrastructure has already been created.

Recent models of the data economy reinforce this cumulative logic. Farboodi and Veldkamp (2026) treat data as an economically valuable intangible asset generated through activity and experimentation. Their framework helps explain why data-intensive production can contribute to divergence in firm size: economic activity creates data, data improve decisions and products, and improved products can generate further activity and additional data. AI can strengthen this feedback mechanism because improvements in prediction and personalization increase the value of well-organized data assets. The relevant advantage is therefore not simply possession of a large dataset but possession of an architecture through which information generated in one part of an organization can be productively redeployed elsewhere.

Evidence at the firm level is consistent with this interpretation. Pastor-Merino et al. (2026) find stronger productivity gains where AI is deployed across multiple business functions rather than confined to isolated applications. A firm that develops a common data and AI layer can use it for demand forecasting, customer service, fraud detection, inventory management, marketing, compliance, and HR decisions. A firm without such an architecture may have to treat each application as a separate project, repeatedly incurring costs for data preparation, integration, contracting, and governance. Scope economies thus reinforce the capability filter mechanism: organizations with mature digital systems face a lower incremental cost of experimenting with each additional AI use. The scope advantage is not fixed, however. Reorganization can make additional activities compatible with a common architecture, consistent with evidence that firms change hierarchical structure and task allocation in response to economic shocks (Garicano, 2000; Caliendo and Rossi-Hansberg, 2012; Caliendo et al., 2015; Brambilla et al., 2025).

The same mechanism may be especially important in government. A high-quality digital identity system is not valuable only for one service. Combined with interoperable registries and secure data-sharing rules, it can support taxation, social protection, licensing, health services, education administration, procurement, and disaster response. This is one reason why digital public infrastructure can generate high social returns: the underlying assets are reusable across agencies and policy domains (Cristia and Vlaicu, 2022; OECD and CAF, 2023). Fragmented government information systems have the opposite effect. If each ministry creates separate identification systems, databases, procurement arrangements, and verification processes, potential scope economies are lost even when each agency is individually digitized.

These dynamics help sharpen the statement that "the data-rich get richer." The advantage of an early mover is not simply a larger stock of information. It is a lower marginal cost of creating the next application because the organization already possesses interoperable data, reusable software, governance systems, skilled personnel, and institutional routines. Foundation models potentially magnify this effect because a common model can be adapted to many tasks. At the same time, the value of such reuse depends on data quality, domain knowledge, evaluation, and governance; a generic model does not eliminate these complementary requirements.

For LAC, this mechanism implies that differences in digital architecture can generate persistent differences in returns even after access to frontier AI models becomes widespread. Firms and governments with integrated systems can add applications sequentially at declining incremental cost, whereas those with siloed or paper-based systems must first pay the fixed costs of creating the underlying digital infrastructure. Economies of scope therefore interact strongly with complementarities and scale: they reward organizations that can build common assets once, reuse them repeatedly, and learn across applications, producing another channel through which initial digital capability gaps can widen over time.

4.5 Labor Market Exposure and Unequal Capacity to Benefit

Finally, uneven returns operate through workers as well as firms. AI affects occupations at the task level, and technological exposure does not map mechanically into either displacement or productivity gains. A task can be automated, complemented, reorganized, or left largely unchanged depending on the technology, the worker's skills, the surrounding workflow, and the employer's capacity to redesign production. Labor market outcomes therefore reflect the same complementarities emphasized elsewhere in the framework. Equation (14) summarizes this distinction: occupational exposure e_o is filtered through task complementarity c_o and the employer's realized capability q_i^*. Workers who have access to appropriate digital

infrastructure, complementary skills, and organizations capable of deploying AI productively may benefit from augmentation, while observationally similar workers in less capable firms may receive little benefit from the same frontier technology. This logic is closely related to the complementarity-adjusted exposure framework of Pizzinelli et al. (2023).

This distinction is particularly important in LAC. The World Bank–ILO analysis finds substantial potential exposure to generative AI but also emphasizes that inadequate digital infrastructure can prevent millions of workers from realizing possible productivity gains (Gmyrek et al., 2024; World Bank and ILO, 2024). Using harmonized household surveys for 14 Latin American countries, Ciaschi et al. (2025) likewise document substantial heterogeneity across worker groups and show that estimated distributional effects change importantly when measures distinguish between potential substitution and task complementarity. Exposure is relatively high among several groups that are already more connected to formal and skill-intensive labor markets, illustrating a central tension: workers most exposed to AI may simultaneously be the workers best positioned to benefit from it.

Evidence outside the region reinforces the importance of distinguishing technological exposure from realized labor market outcomes. Using French firm-level data, Aghion et al. (2025) find that AI-adopting firms are initially larger, more productive, and more skill intensive and experience increases in sales and employment following adoption, on average. Importantly, employment effects differ across types of AI use. This evidence is difficult to reconcile with a simple model in which occupational exposure translates directly into job destruction. Instead, productivity gains can expand firm output and labor demand even while the composition of tasks changes.

Experimental studies reveal a related distributional margin within firms. Generative AI has often produced larger measured productivity improvements for less experienced workers performing particular tasks (Brynjolfsson et al., 2025; Cui et al., 2026). This suggests that AI can sometimes substitute for experience or diffuse knowledge that was previously concentrated among higher-performing workers. Yet this potentially equalizing effect operates only among workers who actually use the technology in an environment where it is productive. Adoption itself is unequal: Humlum and Vestergaard (2025b) document substantial differences in ChatGPT adoption across workers. Unequal access, differences in digital skills, and variation in employer adoption can therefore offset or even reverse within-task equalization.

Early causal evidence also suggests that adjustment may initially occur through task composition rather than aggregate employment. In Denmark, rapid chatbot adoption has been associated with changes in tasks and occupational mobility but little average effect on earn-

ings or hours over the initial period studied (Humlum and Vestergaard, 2025a). This is consistent with the J-curve logic developed above: firms and workers may adopt AI relatively quickly, while larger changes in organizational structure, occupational demand, wages, and employment emerge more slowly.

For developing economies, weak digital infrastructure can consequently provide a temporary "buffer" from automation while limiting access to augmentation. Informality creates an additional divide because informal firms and workers are less likely to operate within standardized, data-rich production systems where AI tools can be readily integrated. Geographic disparities matter for similar reasons: workers in major cities may have better connectivity, employers, training opportunities, and complementary services than those in peripheral regions.

The distribution of AI gains will therefore depend not only on which occupations are technologically exposed but also on which workers and firms possess the complementary assets needed to translate exposure into productivity. This is the labor market counterpart of the capability filter in Section 4.1. Digital transformation can narrow some productivity gaps—for example, by helping less experienced workers perform complex tasks—while widening others through unequal adoption across firms, workers, and locations. Understanding both margins is essential for assessing whether AI ultimately broadens or concentrates economic opportunity in LAC.

5 Policy Implications for LAC

The framework above implies a dual challenge for governments: deepen productive digital adoption and broaden access to the resulting gains. These objectives are related but not identical. In the model, lowering the cost of one input raises returns only when other complements are sufficiently available, while equation (6) implies that firms can remain below the adoption threshold after a partial subsidy. This favors policies that diagnose binding constraints and coordinate complementary investments rather than treating technology acquisition as an objective in itself. The appropriate mix will therefore differ across countries, sectors, and firm types in LAC.

5.1 Building Foundational Layers Strategically

The first priority is to strengthen connectivity, compute access, energy reliability, and basic digital capabilities with explicit attention to complementarities. Connectivity is necessary for

most advanced applications, but its productivity payoff depends on devices, skills, reliable electricity, and useful services. In terms of equation (1), expanding a nonbinding capacity produces no immediate intensive margin gain if another input remains the bottleneck; equation (6) nevertheless shows that lower costs can still raise the extensive margin probability of adoption. Infrastructure policy should therefore be informed by local diagnostics of adoption barriers rather than coverage targets alone.

This does not imply abandoning universal access objectives. Basic connectivity can have high welfare value even where near-term commercial returns are modest, particularly through access to government, health, education, and financial services. High-end compute infrastructure should be assessed differently because data centers, cloud regions, and high-performance computing involve large fixed costs and may be viable only where energy, skills, connectivity, and demand reach sufficient scale. Smaller economies may gain more from coordinated regional procurement or shared cloud capacity than from reproducing the full technology stack domestically.

Human capital is part of the same foundational layer. Recent firm evidence indicates that AI-related productivity gains are larger when adoption is accompanied by software, data, workforce training, or strong managerial capability (Aldasoro et al., 2026; Kazakis, 2026). For LAC, this points toward training linked to concrete workplace applications and managerial decision-making rather than stand-alone digital literacy programs. Bundling technology with implementation support, data management, and process redesign is more likely to move firms across the threshold in equation (6) than subsidizing licenses or equipment alone.

Sequencing also matters. Where basic connectivity and data systems are weak, advanced AI subsidies may simply raise demand for imported tools without generating much organizational learning. Where those foundations are already in place, smaller interventions aimed at management, integration, or specialized skills may have high returns. This suggests using readiness diagnostics to distinguish places where foundational investment should come first from those where the main constraint is adoption capability. The model does not imply a single universal sequence, but it does imply that policy packages should be internally coherent: investments aimed at raising q_i should address the complement that is actually limiting effective capability.

5.2 Strengthening Digital Government

A second priority is to modernize digital government. Governments are not only regulators; they are major producers and users of data-intensive services. They can coordinate infrastructure, establish interoperability and governance standards, and exploit economies of scale and scope. High-quality digital identity, interoperable registries, secure data-sharing protocols, and common payment or notification rails can be reused across taxation, social protection, health, education, licensing, and disaster response. In equation (2), such investments raise Θ_i by increasing the number and value of applications that draw on the same core infrastructure.

The largest gains are unlikely to come from placing existing paper procedures online. They require redesigning roles, information flows, and decision rules around digital capabilities. The permit example in the subsection on workflow redesign illustrates the principle: citizens should not remain couriers between agencies simply because forms have become electronic. Reusable digital public infrastructure can reduce the marginal cost of adding services and improve consistency across agencies (Cristia and Vlaicu, 2022).

AI also raises governance requirements. Agencies need procedures for validating models, assigning responsibility, monitoring errors and distributional effects, protecting sensitive data, and preserving human review. Procurement should therefore evaluate the full service transformation—data readiness, integration, training, and monitoring—rather than the acquisition of an AI tool in isolation.

5.3 Addressing Information Barriers and Coordination Failures

A third priority is to reduce the information barriers that prevent firms, especially SMEs, from identifying productive digital and AI applications. Managers must decide not only whether to adopt a technology but also where to use it, how to reorganize complementary tasks, what data to collect, and how to evaluate performance. These decisions require expertise that many firms do not possess internally, while the returns to experimentation may be only partly appropriable if successful practices are easily observed or imitated by competitors. As a result, firms may underinvest in learning and experimentation even when the social returns to discovering effective applications are substantial. Policy can reduce these information barriers through sector-focused extension services, diagnostic support, demonstration projects, vetted technology providers, and adoption support tied to measurable process changes rather than technology purchases alone.

Coordination failures are distinct. The return to one investment may depend on complementary investments by other actors that are not undertaken simultaneously. For example, an SME may have little incentive to invest in advanced digital systems if local providers of cloud, cybersecurity, or data management services are absent; technology providers may in turn be reluctant to develop specialized services where they expect demand to remain limited. Similar problems arise when firms would benefit from common data standards, shared sectoral platforms, interoperable payment or identification systems, or pools of workers with specialized skills, but no individual firm has sufficient incentive to bear the fixed cost of creating them. In the framework of this paper, such coordination failures can leave several complementary inputs simultaneously below the levels needed to raise effective capability q_i or to cross the adoption threshold in equation (6). Individually rational decisions can therefore produce a collectively low-adoption equilibrium even when coordinated investment would generate positive returns. Equation (11) formalizes this possibility: when the local return shifter $A_{\ell,t}$ rises with the share of adopters, firms' incentives are strategic complements and expectations of low adoption can become self-confirming.

This creates a role for policy that goes beyond providing information or subsidizing individual adopters. Governments and development institutions can help coordinate complementary investments by convening firms and technology providers, establishing common technical and data standards, supporting shared digital infrastructure, organizing sectoral training linked to expected demand, and using procurement or matching programs to create a critical mass of adopters. Such interventions are most compelling when they address identifiable interdependencies rather than simply lowering the private cost of technology. Recent evidence reinforces the importance of this distinction: AI effects vary sharply across tasks and organizational settings (Dell'Acqua et al., 2026b), productivity gains often depend on how technologies are integrated into existing workflows (Cui et al., 2026), and firm-level gains appear stronger when AI is deployed across multiple functions rather than as an isolated application (Pastor-Merino et al., 2026). The objective should therefore be to coordinate the ecosystem of complementary investments needed for productive adoption, not merely to increase the number of firms purchasing digital or AI tools.

5.4 Using Regional Scale and Scope

The scale terms in equations (3) and (7) are particularly relevant for LAC. Many countries are individually too small to support every layer of the AI ecosystem efficiently. Regional cooperation can pool demand for cloud and cybersecurity services, support compatible digital

identity and electronic signature standards, coordinate cross-border payment and data transfer frameworks, and share specialized regulatory expertise. These arrangements allow smaller economies to access capabilities that would be costly to reproduce nationally.

Regional integration can also enlarge the effective market for digital firms. Common technical standards and interoperable systems reduce the fixed cost of serving multiple countries, while fragmented rules can turn each border into a new barrier. The aim need not be complete regulatory uniformity but compatibility where fragmentation prevents scale without delivering commensurate social benefits. For Caribbean economies in particular, pooled procurement and shared infrastructure can help offset the scale disadvantages described in Section 3.

5.5 Managing Labor Market Adjustment

A further challenge is to manage labor market transitions without assuming either mass displacement or frictionless adjustment. Exposure measures identify tasks that AI may affect, but they do not determine whether workers will be displaced, augmented, or reassigned. This distinction is especially important in LAC. The analysis of 14 countries by Ciaschi et al. (2025) finds substantial heterogeneity across worker groups and shows that complementarity-adjusted exposure measures can change the apparent incidence across income and education groups. The World Bank–ILO evidence similarly emphasizes the interaction between occupational exposure and the digital infrastructure needed to realize augmentation gains (Gmyrek et al., 2024).

Policy should therefore emphasize adaptability. Training systems should respond to changing task content and remain closely linked to employers and actual technologies; social protection should facilitate mobility; and employment services can use information on skill demand to guide workers toward occupations in which AI is more complementary. Governments should also avoid basing adjustment policy on technological exposure alone. Early administrative evidence from Denmark finds extensive adoption and task restructuring but little average effect on earnings or hours in the first two years (Humlum and Vestergaard, 2025a). This argues for monitoring realized transitions and targeting support to workers who experience displacement or earnings losses.

Distribution matters on the gains side as well. Several experiments suggest that less experienced workers can obtain larger productivity improvements from AI tools (Brynjolfsson et al., 2025; Cui et al., 2026), potentially compressing some performance gaps. But unequal access and differences in digital skills can push in the opposite direction; ChatGPT adoption

itself varies substantially across worker groups (Humlum and Vestergaard, 2025b). Effective access policies should therefore combine availability with training and organizational changes that enable productive use.

5.6 Promoting Competition, Data Governance, and Evaluation

Finally, strong economies of scale and scope require a competition policy framework that preserves the efficiency gains from large digital networks while limiting practices that allow temporary technological advantages to become durable barriers to entry. Concentration by itself is not necessarily evidence of inefficient market power: large fixed costs, network effects, and the ability to spread investments in data, cybersecurity, and computing over many users can make scale socially productive (Goldfarb and Tucker, 2020). The concern is instead that scale may become self-reinforcing when incumbents combine large user bases with proprietary data and restrict rivals' access to interfaces, complementary services, or customers. Theoretical work shows that data-driven learning can produce indirect network effects under which markets tip toward a dominant firm and innovation incentives subsequently weaken (Prüfer and Schottmüller, 2021). Equations (12)–(13) provide a simple demand-side counterpart in the present framework: adoption lowers marginal cost, expands market share, and passes part of the gain to consumers, while scale and data feedbacks can amplify concentration among early adopters. Competition policy should therefore distinguish between concentration generated by genuine productive efficiencies and conduct that raises artificial entry or switching costs. For LAC, this distinction is particularly important because overly aggressive restrictions on scale could discourage investment in relatively small markets, while weak contestability could allow dominant global or regional platforms to extend existing advantages across adjacent digital markets.

Data governance, portability, and interoperability can help address these problems, but their design requires attention to both competition and incentives to invest. Portability can reduce the cost to users of switching providers, while interoperability can allow competing firms to connect to established networks without reproducing the entire installed base. Economic theory nevertheless indicates that their effects depend on market structure and implementation. These findings caution against treating portability as synonymous with simply transferring a static copy of a user's data. Effective contestability may instead require technically usable standards, timely access, common interfaces, and rules that allow users to move or authorize access to information without losing functionality. At the same time, privacy, cybersecurity, and commercially sensitive information place legitimate limits on data

sharing. For LAC, common or compatible regional standards could be particularly valuable because firms otherwise face different technical and regulatory requirements in relatively small national markets. The objective should therefore be to lower unnecessary switching and entry barriers while retaining incentives to invest in data creation, infrastructure, and service quality.[7]

Because the relevant technologies and business models are evolving rapidly, these competition and governance policies should be accompanied by systematic evaluation. The framework developed in this paper implies that observing little response to an intervention does not necessarily mean that the underlying technology has low returns. Equation (6) shows that a subsidy or infrastructure improvement may have little measured effect if another complement remains binding; conversely, broad subsidies may largely accrue to firms that would have adopted without support. Recent empirical research further demonstrates that AI effects differ substantially across tasks, workers, firms, and organizational settings (Dell'Acqua et al., 2026b; Cui et al., 2026; Aldasoro et al., 2026; Pastor-Merino et al., 2026). Policy evaluation should therefore measure not only technology adoption but also changes in complementary investment, workflow, productivity, entry, prices, employment, and distributional outcomes. Where feasible, pilots, phased implementation, and experimental or quasi-experimental designs can help distinguish whether weak outcomes reflect an ineffective intervention, an unresolved complementary constraint, or simply insufficient time for organizational adjustment. An adaptive approach is especially appropriate for AI policy: regulations and support programs should be designed so that evidence from implementation can inform their subsequent scale, targeting, and design.

6 Conclusion

Digitalization and AI constitute a promising new wave of GPTs, but technological capability does not translate mechanically into economic value. This paper emphasizes four connected properties that shape realized returns: complementarities among technology, data, skills, and organizational capital; the need for workflow and business process redesign; economies of scale; and economies of scope. The model in Section 3.1 brings these mechanisms together while allowing AI to substitute for some task-specific skills. Equation (1) separates that substitution margin from genuine bottlenecks in data, implementation skills, and organizational readiness;

[7]For example, Jeon et al. (2023) show that interoperability can increase consumer surplus when switching costs are high, while the effects of data portability depend on how firms adjust compatibility and pricing decisions.

equation (6) shows why adoption can remain discontinuous, and equations (10) and (11) show how adoption advantages can become cumulative within firms and self-reinforcing across local ecosystems. The demand and labor market extensions in equations (12)–(14) connect these firm-level mechanisms to competition, consumer pass-through, and heterogeneous worker outcomes. The implication is that uneven returns are not an incidental feature of digital transformation. They follow naturally from the economics of technologies whose productivity depends on complementary investments, substantial fixed adjustment costs, and the capacity to reuse digital assets across users and applications.

This framework also helps reconcile the rapid diffusion of generative AI tools with the mixed evidence on broader productivity effects. Experimental studies document sizable gains for particular tasks and workers, but these gains vary with experience, task characteristics, and how AI is embedded in work processes (Brynjolfsson et al., 2025; Dell'Acqua et al., 2026b; Cui et al., 2026). Emerging firm-level evidence likewise indicates that productivity gains are stronger when AI adoption is accompanied by complementary software, data, skills, management capabilities, and deployment across multiple business functions (Aldasoro et al., 2026; Kazakis, 2026; Pastor-Merino et al., 2026). At the same time, broader workplace evidence suggests that organization-wide adjustment can proceed considerably more slowly than tool adoption itself (Dillon et al., 2026; Humlum and Vestergaard, 2025a). The relevant horizon is therefore longer than the software adoption cycle. The productivity effects of AI will depend on experimentation, organizational learning, investment in intangible capital, and the redesign of production processes, much as earlier GPTs produced their largest aggregate gains only after complementary innovations had diffused (Brynjolfsson et al., 2021).

For LAC, these mechanisms create both a significant opportunity and a risk of further divergence. Technologies that lower the cost of prediction, communication, coordination, and knowledge-intensive services can help firms overcome some traditional barriers associated with distance and market size and can improve the quality and efficiency of public services. Yet the capacity to capture these benefits is distributed unevenly. Differences in connectivity, managerial capability, workforce skills, firm size, access to finance, data infrastructure, and state capacity can concentrate gains among already-advantaged firms, workers, sectors, and locations. The same asymmetry appears in labor markets. Regional evidence shows substantial heterogeneity in exposure to AI and demonstrates that distributional conclusions depend importantly on whether technology substitutes for worker tasks or complements them (Gmyrek et al., 2024; Ciaschi et al., 2025). Moreover, limited digital infrastructure may temporarily insulate some workers from automation while preventing them from benefiting from AI-enabled productivity improvements. Access to frontier technologies alone is therefore

an incomplete development strategy; what matters is whether firms, workers, and governments possess the complementary capabilities needed to use them productively.

The policy challenge is consequently not simply to accelerate AI adoption, nor to attempt to equalize adoption across all firms and locations. It is to identify the constraints that prevent potentially productive users from reaching the capability thresholds described in the paper while ensuring that the gains from scale and scope do not unnecessarily restrict future entry or access. This requires a combination of foundational digital infrastructure, skills and managerial capabilities, interoperable digital public infrastructure, support for experimentation and organizational change, regional approaches where national scale is insufficient, and competition and data governance institutions that preserve contestability. Because the technology and the evidence are evolving rapidly, policy should also be adaptive: major interventions should incorporate measurement and evaluation from the outset and be revised as evidence accumulates about productivity, employment, competition, and distributional effects. Whether digitalization and AI narrow or widen LAC's long-standing productivity and inequality gaps will ultimately depend less on access to the technology itself than on the region's capacity to build, coordinate, and broadly diffuse the complementary capabilities that turn technological potential into sustained economic value.

References

Acemoglu, D. (2025). The Simple Macroeconomics of AI. *Economic Policy*, 40(121):13–58.

Aghion, P., Bunel, S., Jaravel, X., Mikaelsen, T., Roulet, A., and Søgaard, J. (2025). How Different Uses of AI Shape Labor Demand: Evidence from France. *AEA Papers and Proceedings*, 115:62–67.

Agrawal, A., Gans, J., and Goldfarb, A. (2018). *Prediction Machines: The Simple Economics of Artificial Intelligence*. Harvard Business Review Press.

Agrawal, A., Gans, J., and Goldfarb, A. (2022). *Power and Prediction: The Disruptive Economics of Artificial Intelligence*. Harvard Business Review Press.

Aldasoro, I., Gambacorta, L., Pal, R., Revoltella, D., Weiss, C., and Wolski, M. (2026). AI Adoption, Productivity and Employment: Evidence from European Firms. EIB Working Paper 2026/02, European Investment Bank.

Brambilla, I., César, A., Falcone, G., and Porto, G. (2025). Organizational Hierarchies and Export Destinations. *The World Bank Economic Review*, 39(2):229–259.

Bresnahan, T. F. and Trajtenberg, M. (1995). General Purpose Technologies: 'Engines of Growth'? *Journal of Econometrics*, 65(1):83–108.

Brynjolfsson, E., Li, D., and Raymond, L. R. (2025). Generative AI at Work. *Quarterly Journal of Economics*, 140(2):889–942.

Brynjolfsson, E., Rock, D., and Syverson, C. (2021). The Productivity J-Curve: How Intangibles Complement General Purpose Technologies. *American Economic Journal: Macroeconomics*, 13(1):333–372.

Caliendo, L., Monte, F., and Rossi-Hansberg, E. (2015). The Anatomy of French Production Hierarchies. *Journal of Political Economy*, 123(4):809–852.

Caliendo, L. and Rossi-Hansberg, E. (2012). The Impact of Trade on Organization and Productivity. *Quarterly Journal of Economics*, 127(3):1393–1467.

Ciaschi, M., Falcone, G., Garganta, S., Gasparini, L., Bertín, O., and Ramírez-Leira, L. (2025). The Potential Distributive Impact of AI-Driven Labor Changes in Latin America. Working paper, Inter-American Development Bank.

Cristia, J. P. and Vlaicu, R. (2022). *Digitalizing Public Services: Opportunities for Latin America and the Caribbean*. Inter-American Development Bank.

Cui, K. Z., Demirer, M., Jaffe, S., Musolff, L., Peng, S., and Salz, T. (2026). The Effects of Generative AI on High-Skilled Work: Evidence from Three Field Experiments with Software Developers. *Management Science*. Published online February 27, 2026.

Dell'Acqua, F., Ayoubi, C., Lifshitz, H., Sadun, R., Mollick, E., Mollick, L., Han, Y., Goldman, J., Nair, H., Taub, S., and Lakhani, K. R. (2026a). The Cybernetic Teammate: A Field Experiment on Generative AI and Teamwork. *Organization Science*, 37(4):1217–1242.

Dell'Acqua, F., McFowland, III, E., Mollick, E., Lifshitz, H., Kellogg, K. C., Rajendran, S., Krayer, L., Candelon, F., and Lakhani, K. R. (2026b). Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of Artificial Intelligence on Knowledge Worker Productivity and Quality. *Organization Science*, 37(2):403–423.

Dillon, E. W., Jaffe, S., Immorlica, N., and Stanton, C. T. (2026). Shifting Work Patterns with Generative AI. *American Economic Review: Insights*. Forthcoming.

Dixit, A. K. and Stiglitz, J. E. (1977). Monopolistic Competition and Optimum Product Diversity. *American Economic Review*, 67(3):297–308.

ECLAC (2025). Latin America and the Caribbean Accelerate the Adoption of Artificial Intelligence, Though Challenges Remain in Investment, Talent and Governance. Technical report, ECLAC.

Farboodi, M. and Veldkamp, L. (2026). A Model of the Data Economy. *The Review of Economic Studies*. Corrected proof, article rdag074.

Gambacorta, L., Qiu, H., Shan, S., and Rees, D. M. (2026). Generative AI and Labour Productivity: A Quasi Experiment on Coding. *Journal of Financial Stability*, 84:101543.

Garicano, L. (2000). Hierarchies and the Organization of Knowledge in Production. *Journal of Political Economy*, 108(5):874–904.

Gmyrek, P., Winkler, H., and Garganta, S. (2024). Buffer or Bottleneck? Employment Exposure to Generative AI and the Digital Divide in Latin America. ILO Working Paper 121, International Labour Organization and World Bank.

Goldfarb, A. and Tucker, C. (2020). Digital Economics. *Journal of Economic Literature*, 58(1):3–43.

GSMA (2023). The Mobile Economy Latin America 2023. Technical report, GSM Association.

Hausmann, R., Rodrik, D., and Velasco, A. (2008). Growth Diagnostics. In Serra, N. and Stiglitz, J. E., editors, *The Washington Consensus Reconsidered: Towards a New Global Governance*, pages 324–355. Oxford University Press.

Helpman, E., editor (1998). *General Purpose Technologies and Economic Growth*. MIT Press.

Humlum, A. and Vestergaard, E. (2025a). Still Waters, Rapid Currents: Early Labor Market Transformation under Generative AI. Working Paper 33777, National Bureau of Economic Research. Revised March 2026.

Humlum, A. and Vestergaard, E. (2025b). The Unequal Adoption of ChatGPT Exacerbates Existing Inequalities among Workers. *Proceedings of the National Academy of Sciences*, 122(1):e2414972121.

IDB (2021). Artificial Intelligence and the Future of Work in Latin America and the Caribbean: A Framework for Action. Technical report, Inter-American Development Bank.

IDB (2022). The 360 on Digital Transformation in Firms in Latin America and the Caribbean. Technical report, Inter-American Development Bank.

ITU (2024). Facts and Figures 2024: Internet Use. Technical report, International Telecommunication Union.

Jeon, D.-S., Menicucci, D., and Nasr, N. (2023). Compatibility Choices, Switching Costs, and Data Portability. *American Economic Journal: Microeconomics*, 15(1):30–73.

Jones, C. I. and Tonetti, C. (2020). Nonrivalry and the Economics of Data. *American Economic Review*, 110(9):2819–2858.

Jovanovic, B. and Stolyarov, D. (2000). Optimal Adoption of Complementary Technologies. *American Economic Review*, 90(1):15–29.

Kazakis, P. (2026). Conditional Gains: When AI Investment Enhances Firm Efficiency. *Scottish Journal of Political Economy*, 73(4):e70055.

Mandon, P. (2025). Beyond the AI Divide: A Straightforward Approach to Identifying Global and Local Overperformers in AI Preparedness. *Digital Business*, 5(2):100136.

McElheran, K., Yang, M.-J., Kroff, Z., and Brynjolfsson, E. (2026). The Adoption of Industrial AI in America. *AEA Papers and Proceedings*, 116:20–25.

Noy, S. and Zhang, W. (2023). Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence. *Science*, 381(6654):187–192.

OECD (2021). *The Digital Transformation of SMEs*. OECD Publishing.

OECD (2023). Competition Policy in the Digital Age. Technical report, OECD Publishing.

OECD (2024). AI in Society 2024. Technical report, OECD Publishing.

OECD (2025). *Results from TALIS 2024: The State of Teaching*. TALIS. OECD Publishing, Paris.

OECD and CAF (2023). *Digital Government Review of Latin America and the Caribbean*. OECD Publishing.

Oxford Insights (2024). Government AI Readiness Index 2024. Technical report, Oxford Insights.

Pastor-Merino, A., Martínez-Barbero, X., Vicente, M. R., and Domenech, J. (2026). Does AI Boost Firm Productivity? A Web Scraping and LLMs Approach. *Telecommunications Policy*, 50(2):103138.

Pisu, M. et al. (2021). Digitalisation and Productivity: A Story of Complementarities. *OECD Economic Outlook*, 2021(1):55–83.

Pizzinelli, C., Panton, A. J., Tavares, M. M., Cazzaniga, M., and Li, L. (2023). Labor Market Exposure to AI: Cross-Country Differences and Distributional Implications. IMF Working Paper 2023/216, International Monetary Fund.

Prüfer, J. and Schottmüller, C. (2021). Competing with Big Data. *The Journal of Industrial Economics*, 69(4):967–1008.

Stanford Institute for Human-Centered Artificial Intelligence (2025). The 2025 AI Index Report. Technical report, Stanford University.

UNCTAD (2021). Estimates of Global E-Commerce 2019 and Preliminary Assessment of COVID-19 Impact on Online Retail 2020. UNCTAD Technical Notes on ICT for Development 18, United Nations.

World Bank (2021). *World Development Report 2021: Data for Better Lives*. World Bank.

World Bank and ILO (2024). Quantifying the Jobs Potential of AI in Latin America and the Caribbean. Results brief, World Bank.