

# The Value of Organizational Learning Technologies\*

**Martin Beraja**

UC Berkeley

**Eduard Talamàs**

IESE Barcelona

## Abstract

Organizations learn over time. They build organizational capital and form beliefs about their fundamentals. We conceptualize AI as an organizational learning technology that promises to accelerate both forms of learning. We show that, in a large class of firm dynamics models, the value of organizational learning technologies (VOLT) is governed by two simple statistics: the relative size and lifespan of mature firms. In the United States, VOLT is on the order of one GDP — implying that accelerating organizational learning could double aggregate output. Across industries, VOLT exhibits substantial variation and is orthogonal to existing measures of AI exposure based on adoption feasibility, complementing them in assessing AI’s potential impact. Much of VOLT reflects increases in average firm lifespans rather than productivity, revealing firm longevity as a powerful channel through which organizational learning technologies can affect output.

## 1 Introduction

Organizations learn over time. They build organizational capital, becoming more productive with age (Prescott and Visscher, 1980; Atkeson and Kehoe, 2005; Eifeldt and Panikolaou, 2013), and form beliefs about their future fundamentals, which shape their entry and exit decisions (Jovanovic, 1982; Pakes and Ericson, 1998). By shaping productivity dynamics and firm selection, frictions in the process of organizational learning depress aggregate productivity (Hsieh and Klenow, 2014) and hinder the reallocation of resources across firms (Buera and Shin, 2013; Midrigan and Xu, 2014).

---

\*Date Printed: February 6, 2026.

Martin Beraja: mberaja@berkeley.edu. Eduard Talamàs: etalamas@iese.edu. We thank Karan Makkar for excellent research assistance. Eduard acknowledges the financial support of IESE through the High Impact Initiative-course 2024/2026. All errors are our own.

Recent advances in Artificial Intelligence (AI) promise to accelerate organizational learning. Large Language Models (LLMs) are primarily being used as learning tools (Appel et al., 2025; Chatterji et al., 2025; Tomlinson et al., 2025), while AI-driven organizational memory systems — such as *Glean* and *o9's Digital Brain* — enable firms to build organizational capital faster by aggregating, indexing, and making retrievable information that is otherwise dispersed across the organization. In parallel, AI prediction systems — such as *Expected Parrot*, *Kruncher* and *Shopify's SimGym* — help forecast fundamentals like future demand and costs, improving assessments about the viability of ideas before they are launched. Indeed, AI agents already outperform experienced human venture capital analysts in predicting early-stage startup survival (Vismara et al., 2025).

In this paper, we conceptualize technologies that accelerate organizational learning and show that their value is governed by two simple statistics in a large class of models of firm dynamics. We study a technology that accelerates learning early in a firm's life: average productivity when young increases to its expected mature productivity — capturing faster organizational capital accumulation — and firms that would otherwise exit while young are screened out before entry — capturing faster learning about fundamentals. We define the Value of Organizational Learning Technologies (VOLT) as the resulting increase in aggregate output.

Specifically, we consider a class of models of firm dynamics à la Hopenhayn (1992) with organizational learning. We define VOLT as the ratio of stationary-equilibrium output in the accelerated-learning counterfactual economy to that in the factual economy, and we show that it is determined by two statistics: the ratio of mature-to-average firm size  $y$ , and the ratio of mature-to-average firm lifespan  $\ell$ . Formally,

$$\text{VOLT} = (y \cdot \ell)^{1-\nu} \tag{1.1}$$

where  $\nu$  is the span of control parameter (Lucas, 1978) that governs decreasing returns to scale at the firm level. The relative firm size  $y$  is larger when mature firms have accumulated more organizational capital and are more positively selected on productivity. The relative firm lifespan  $\ell$  is larger when exit rates among young firms are higher.

Using 2023 U.S. establishment-level data from the Business Dynamics Statistics (BDS), we compute VOLT under alternative definitions of maturity. As one benchmark, we define mature establishments as those in the left-censored age bin — i.e., establishments born before 1977 and thus older than 46 years in 2023. These establishments are on average about 3 times larger than the average establishment ( $y \approx 3$ ), and they exhibit much lower exit rates: their implied expected lifespan is 68 years versus an average establish-

ment lifespan of 10 years ( $\ell \approx 7$ ). Taking these statistics to the VOLT formula (1.1) with a standard value for the span-of-control parameter  $\nu = 0.75$  (Hopenhayn, 2014) yields

$$\text{VOLT} \approx (3 \cdot 7)^{1-0.75} \approx 2.$$

In other words, accelerating organizational learning has the potential to double U.S. GDP. Using shorter horizons for the learning acceleration yield smaller but still sizable effects: taking mature firms to be those older than 20 years implies a 60% increase in stationary equilibrium output, while using a maturity threshold of 5 years implies a 22% increase in output.

Approximately one quarter of the increase in output is driven by higher Total Factor Productivity (TFP), while the remaining reflects a larger mass of firms, as faster learning increases the average firm lifespan in the economy.<sup>1</sup> This finding indicates that greater firm longevity is a powerful mechanism through which organizational learning can affect aggregate output in the class of models we consider, beyond increases in firm productivity which have been the focus of much of the literature.

Disaggregating the VOLT across sectors reveals substantial heterogeneity. We find that VOLTs are largely orthogonal to standard measures of LLM exposure (Eloundou et al., 2024) revealing a rich off-diagonal pattern: sectors with relatively low exposure often harbor large gains from faster learning, while sectors where LLMs appear relatively easy to adopt often yield relatively modest gains.

This paper contributes to several strands of research: the assessment of AI’s transformative potential, the investigation of firm dynamics with organizational learning, and the measurement of AI’s organizational impact.

*AI’s Transformational Potential.* Current debates about AI’s transformative potential (e.g., Agrawal et al., 2025) have so far focused on automation of production (e.g., Acemoglu, 2025; Ide and Talamàs, 2025; Ide, 2025; Beraja and Zorzi, 2025) — finding modest aggregate productivity gains — or the extent to which AI’s potential to accelerate scientific discovery (e.g., Amodei, 2024; Mullainathan and Rambachan, 2025) can lead to explosive growth (e.g., Aghion et al., 2017; Trammell and Korinek, 2025; Jones, 2025; Restrepo, 2025; Davidson et al., 2026; Jones and Tonetti, 2026; Jones, 2026).

In this paper, we abstract from these two important channels and instead conceptualize AI as an organizational learning technology. We show that this complementary perspective has the advantage that the key statistics required to quantify AI’s potential

---

<sup>1</sup> The number of firms is important for aggregate output due to decreasing returns to scale at the firm level (Hopenhayn, 2014).

impact are readily available in existing firm-level data. The implied aggregate effects are substantial, suggesting that AI's macroeconomic impact may be large even if it only increases the speed at which existing organizational learning occurs.

*Firm Dynamics and Organizational Learning.* A foundational literature views the process of organizational learning as an important driver of firm dynamics and heterogeneity (e.g., Jovanovic, 1982; Pakes and Ericson, 1998; Atkeson and Kehoe, 2005; Hsieh and Klenow, 2014; Perla and Tonetti, 2014; Arkolakis et al., 2018). Our measurement exercise is most closely related to Atkeson and Kehoe (2005) and Hsieh and Klenow (2014).

Atkeson and Kehoe (2005) build a model of plant life cycles à la Lucas (1978) and Hopenhayn (1992), and use establishment-level manufacturing data to infer the payments to organization capital. They estimate that payments to organization capital are substantial: roughly 4 percent of manufacturing value added and approximately one-third of the net payments to physical capital. Hsieh and Klenow (2014) show that replacing the U.S. plant life cycle with the Indian or Mexican one lowers aggregate manufacturing TFP by roughly 15–25 percent, underscoring the importance of organizational capital accumulation for aggregate productivity.

In this paper we ask a different question. Rather than measuring the rents accruing to organizational capital, or asking how aggregate outcomes would change if firms accumulated more organizational capital over their lifecycle, we study how aggregate output responds when a new technology — such as AI — accelerates the process of organizational learning. In particular, in contrast to Hsieh and Klenow (2014), we do not ask what would happen if firms learned *more*, but what would happen if they learned faster *what they are already learning*.

*Measurement of AI's Organizational Impact.* A growing body of empirical work quantifies AI's ability to increase productivity and accelerate organizational learning. For example, Brynjolfsson et al. (2025) find that customer-support agents equipped with generative AI assistants increased productivity by 14% on average, with the largest gains accruing to less-experienced workers. This leveling up effect is in line with Noy and Zhang (2023), who demonstrate that in mid-level professional writing tasks, generative AI reduces task completion time by 40% while simultaneously increasing output quality by 18%. These gains are concentrated among lower-performing workers, leading to a substantial compression of the performance distribution. Similarly, in the context of high-end knowledge work, Dell'Acqua et al. (2023) show that management consultants in the bottom half of the skills distribution benefited the most from AI, achieving a 43% performance increase compared to 17% for top performers.

Recent evidence also suggests that AI can significantly alter collaboration inside organizations. In particular, [Dell’Acqua et al. \(2025\)](#) find that AI allows individuals to match the performance of human teams, effectively democratizing access to domain expertise. They show that AI breaks down functional silos — allowing employees outside their core roles to bridge knowledge gaps and produce balanced, holistic solutions — thereby substituting for the complex human coordination traditionally required to accumulate organizational capital.

We complement this research by measuring AI’s aggregate potential for economic impact as an organizational learning technology. Rather than asking how much existing AI tools improve performance in different organizational settings, we ask a complementary question: How much can organizational learning technologies increase aggregate output, and how do the resulting gains vary across sectors? Our results highlight that AI’s impact can arise not only through within-organization productivity gains emphasized in this literature, but also through firm life-cycle dynamics that affect the equilibrium mass of active firms.

## 2 Model

In this section, we lay out a class of models of firm dynamics that we use to measure the value of organizational learning technologies (VOLT). The models in this class are similar to the canonical framework of [Hopenhayn \(1992\)](#) — where firm heterogeneity, entry, and exit are central — but also incorporate organizational learning. [Section 5](#) discusses several extensions of this class of models.

### 2.1 Firm Types

Time is discrete, indexed by  $t \in \{0, 1, \dots\}$ . A firm’s type is a deterministic sequence  $S = (S_0, S_1, \dots)$ , where  $S_k$  is the firm’s state at age  $k$ . The firm does not know its type, but it perfectly learns its state  $S_k$  upon reaching age  $k$ . The state vector can be decomposed as

$$S_k = (z_k, \omega_k) \in \mathcal{Z} \times \Omega.$$

Here,  $z_k \in \mathcal{Z}$  is the firm’s productivity at age  $k$ , while  $\omega_k \in \Omega$  is an informational state that summarizes all relevant history up to age  $k$ . That is,  $\omega_k$  captures the firm’s beliefs about its underlying type: the future sequence of states it will experience.

**Assumption 1. (Rational Expectations).** *Firms subjective belief about any future outcome,*

conditional on their current information  $\omega_k$ , coincides with the true conditional probability.

Following the reduced-form approach in [Atkeson and Kehoe \(2005\)](#), we treat the state paths  $S$  as capturing all the firm-specific factors that shape productivity and survival over a firm's life cycle. In particular, following the literature (e.g., [Prescott and Visscher, 1980](#); [Nonaka, 1994](#); [Grant, 1996](#); [Atkeson and Kehoe, 2005](#); [Argote, 2012](#); [Eisfeldt and Papanikolaou, 2013](#)), we refer to the collection of internally accumulated, non-tradable assets driving firm productivity  $z_k$  as *organizational capital*.

We model firm types as deterministic sequences — albeit unknown to the firm — rather than the more standard approach of specifying a stochastic productivity process and a particular model of belief formation. This formulation has three advantages. First, it is general: our analysis applies to a wide range of environments with varying learning and productivity dynamics. Second, it provides a natural way to conceptualize changes in organizational learning technologies: a new technology alters the sequence of states firms traverse as they age, or the information revealed along those sequences. Third, it clarifies which are the key assumptions behind our main sufficient statistic result.

## 2.2 Firm Production and Exit Choices

Each active firm  $i$  operates with a firm-specific productivity  $z_i$  and labor  $n_i$ . The production technology exhibits decreasing returns to scale:<sup>2</sup>

$$y_i = z_i n_i^\nu, \quad \text{with } \nu \in (0, 1).$$

Following [Lucas \(1978\)](#), we interpret  $\nu$  as a span of control parameter.<sup>3</sup> As in classical models of firm dynamics and organizational capital ([Hopenhayn, 1992](#); [Atkeson and Kehoe, 2005](#)), labor is freely mobile across firms in each period and the decision of how much labor  $n_i$  to hire is static. In [Section 5.2](#), we consider deviations from perfect mobility of labor and static optimization that may arise from, for example, adjustment costs, learning by doing, or borrowing constraints.

The firm observes its productivity  $z_i$  before choosing how much labor  $n_i$  to hire, and

---

<sup>2</sup> This specification is a reduced-form representation of a more general production function  $\hat{z}_i F(n_i, K_i)$ , where  $K_i$  represents other inputs such as capital and intermediates and  $F$  is homogeneous of degree  $\nu$ . If the ratio of these inputs to labor is constant (e.g., due to constant relative factor prices), they can be absorbed into an effective productivity term  $z_i = \hat{z}_i F(1, K_i/n_i)$  so that  $y_i = z_i n_i^\nu$ .

<sup>3</sup> Alternatively, the decreasing returns at the firm level could originate in the demand side of the market, as in [Dixit and Stiglitz \(1977\)](#) and [Melitz \(2003\)](#) (see [Hopenhayn, 2014](#)).

its static problem is maximizing variable profits

$$\pi_i = pz_i n_i^\nu - n_i,$$

given output price  $p$ , where we normalize the wage to 1. The firm's optimal labor, output and variable profits are thus:

$$n^*(z_i) = (pvz_i)^{\frac{1}{1-\nu}}, y^*(z_i) = \frac{1}{p\nu} n_i^*(z) \text{ and } \pi^*(z_i) = \frac{1-\nu}{\nu} n_i^*(z), \quad (2.1)$$

which are proportional to its effective productivity  $z_i^{\frac{1}{1-\nu}}$ . For brevity, from now on we refer to  $z_i^{\frac{1}{1-\nu}}$  simply as firm  $i$ 's *productivity*.

At the end of each period, after signals have been realized and production has occurred, each firm decides whether to continue to the next period or exit. Let  $V(S_k)$  denote the value of a firm at the beginning of age  $k$  given its state  $S_k$ . Each period, the firm earns current profits and then chooses whether to continue operating or exit. The firm's value therefore satisfies:

$$V(S_k) = \pi^*(z_k) - f + \beta \max(0, \mathbb{E}[V(S_{k+1}) | S_k]), \quad (2.2)$$


where  $f$  denotes the fixed cost of operation, the expectation uses the firm's beliefs about future productivities conditional on its state  $S_k$ , and  $\beta < 1$  is the discount factor. Each firm exits the first time the expected continuation value (the term inside the expectation in Equation 2.2) falls strictly below the outside option (normalized to zero). This induces a realized lifespan  $\ell(S)$  for every underlying type  $S$ . To ensure a well-defined stationary distribution, we assume that the lifespan  $\ell(S)$  is finite for all  $S \in \mathcal{S}$ . In other words, all firms exit eventually. Figure 1 illustrates possible productivity sequences of firms that exit at different ages.

To ensure the existence of a balanced growth path — where exit rates are stable and the firm size distribution is stationary as aggregate productivity increases — we assume that the fixed operating cost  $f$  is proportional to profits, so (net) profits grow at the same rate as aggregate productivity (Luttmer, 2007).

**Assumption 2. (Fixed Cost Is Proportional to Mature Profits).** *The fixed cost  $f$  satisfies*

$$f = \alpha \Pi_{>T} \text{ where } \Pi_{>T} \equiv \mathbb{E}_\mu[\bar{\pi}_{>T}(S) | S \in \mathcal{S}_{>T}] \text{ for some } T \geq \tau. \quad (2.3)$$

An assumption like this is needed because we have assumed a constant flow of en-



**Figure 1:** Two possible productivity sequences. The solid lines correspond to ages in which the firm is active. The gray dashed lines illustrate the productivities that the firms would have had beyond the age at which they exit.

trants. As we explain in Section 3.2, such constant flow together with the normalization of fixed costs (Equation 2.3) is convenient for our counterfactual experiments. In particular, it ensures that exit decisions of mature firms are independent on how young organizations learn. In Section 5, we consider an alternative specification where the number of firms is fixed and free entry determines the fixed costs, as in [Atkeson and Kehoe \(2005\)](#).

### 2.3 Organizational Learning

We introduce an age threshold  $\tau$  to partition firms into “young” (age  $k \leq \tau$ ) and “mature” (age  $k > \tau$ ). We also partition the space of all types  $\mathcal{S}$  into two disjoint sets:

- *Exiters*  $\mathcal{S}_{\leq \tau} := \{S \in \mathcal{S} \mid \ell(S) \leq \tau\}$ : The set of types that exit while young.
- *Survivors*  $\mathcal{S}_{> \tau} := \{S \in \mathcal{S} \mid \ell(S) > \tau\}$ : The set of types that survive to mature age.

We consider two forms of organizational learning. The first captures firm learning in the form of organizational capital accumulation that may lead to increases in firm-specific productivity over time. The second form of organizational learning relates to future fundamentals: firms update their beliefs over time ([Jovanovic, 1982](#)), which drives their entry and exit decisions and determines their lifespan  $\ell(S)$ .

We abstract from the specific mechanics of belief formation and make the following assumptions directly about the distribution  $\mu$  of firm types over the set  $\mathcal{S}$  of all possible firm types. Potential entrants can partition the set  $\mathcal{S}$  of all possible types into a collection of subsets. Upon entry, an entrant draws a type according to  $\mu$  from the subset in this collection with the highest expected value. We interpret entry as resulting from effort devoted to idea generation. Choosing a subset corresponds to directing this effort toward generating ideas from that subset.

For now, we assume that potential entrants can identify only a single subset — namely, the full set  $\mathcal{S}$  — and that the expected present value of drawing a type from this set is positive. Under this assumption, entry is optimal. We also assume that survivors are positively selected: firms that survive to maturity are, on average, both more productive and more valuable than those that exit earlier.<sup>4</sup> To state this assumption formally, let the expected mature productivity of type  $S$  at the time of reaching maturity be:

$$\zeta_\tau(S) := \mathbb{E}_\mu \left[ \frac{1}{\ell^*(X) - \tau} \sum_{k=\tau+1}^{\ell^*(X)} z_k(X)^{\frac{1}{1-\nu}} \mid X_\tau = S_\tau \right] \quad (2.4)$$

where  $\ell^*(X)$  denotes the first age  $k > \tau$  at which firm  $X$ 's expected value of continuing falls below the outside option.<sup>5</sup>

**Assumption 3. (Survivors Are Positively Selected).** *Survivors are positively selected on both their expected mature productivity  $\mathbb{E}_\mu[\zeta_\tau(S) \mid S \in \mathcal{S}_{>\tau}] \geq \mathbb{E}_\mu[\zeta_\tau(S) \mid S \in \mathcal{S}_{\leq\tau}]$  and their expected value upon reaching maturity  $\mathbb{E}_\mu[V(S_{\tau+1}) \mid S \in \mathcal{S}_{>\tau}] \geq \mathbb{E}_\mu[V(S_{\tau+1}) \mid S \in \mathcal{S}_{\leq\tau}]$ .*

This assumption captures the idea that exit while young disproportionately eliminates firms with weak organizational capital or poor growth prospects and is consistent with the empirical regularity that mature firms are substantially larger and more productive than young firms (Hsieh and Klenow, 2014).

## 2.4 Stationary Equilibrium

A continuum of firms enters the economy in each period, with the mass of entrants normalized to one. We interpret this as a fixed exogenous flow of new ideas generated by potential firm owners in each period. In Section 5.1, we consider an alternative specifi-

<sup>4</sup> A higher productivity of survivors does not imply a higher value, nor vice versa, because value depends not only on expected average productivity but also on expected lifespan.

<sup>5</sup> Note that  $\ell^*(X)$  coincides with  $\ell(X)$  for survivors. For exitors,  $\ell^*(X)$  is what their lifetime would be if they did not leave before age  $\tau$ .

cation in which the flow of entrants — equivalently, the arrival rate of new ideas — is endogenous.

In a stationary equilibrium, the cross-sectional distribution of firms at any point in time mirrors the dynamic evolution of a single cohort.<sup>6</sup> In particular, the mass  $M$  of active firms in the stationary equilibrium equals firms' expected lifespan  $\bar{\ell}$ :

$$M = \mathbb{E}_\mu [\ell(S)] \equiv \bar{\ell}.$$

Moreover, the cross-sectional aggregate of any flow variable  $x$  — for example, output or labor — is equal to the expected lifetime flow of the representative cohort:

$$X = \mathbb{E}_\mu \left[ \sum_{k=0}^{\ell(S)-1} x_k(S) \right].$$

To clear the goods market, we take the aggregate (demand) expenditure to be fixed at a given level  $E$ :

$$pY = E. \tag{2.5}$$

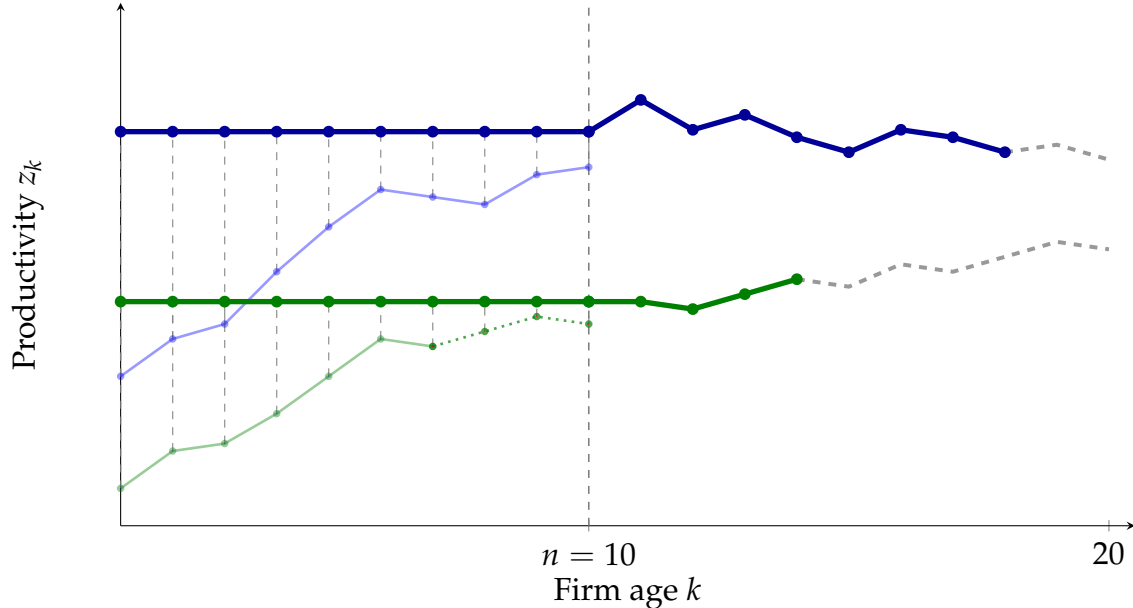
This condition is consistent with two different interpretations of our model. First, we can interpret it as that of an aggregate economy where the labor supply is constant, and thus aggregate expenditure  $pY$  must be fixed. Indeed, in this class of models, labor income is a constant fraction of output,  $N = \nu pY$ , so a fixed labor supply  $N$  implies a fixed expenditure  $pY$ . Second, we can interpret our model as that of a small open economy or industry facing a demand function with unit elasticity, for example as when aggregate output is a Cobb-Douglas aggregator across industries (see Section 4.3). In this case, the labor that is allocated to each industry is also constant.<sup>7</sup>

### 3 VOLT: A Sufficient Statistic Result

We now describe our counterfactual experiment where a technology accelerates organizational learning, and derive a sufficient statistic result for the value of organizational learning technologies (VOLT).

<sup>6</sup> This result is a direct application of the stationary population (Palm) identity.

<sup>7</sup> In Section 5, we consider an extension in which aggregate demand is given by  $p^\theta Y = E$ , where  $\theta$  can capture different elasticities of labor supply (in the aggregate economy) or output demand (in the case of a small open economy or an industry).



**Figure 2:** Possible counterfactual productivity paths for a survivor and an exitor firm type when the maturity threshold is  $\tau = 10$ .

### 3.1 Accelerating Organizational Learning

We model a technology’s effect on organizational learning through two channels: (i) faster accumulation of organizational capital, and (ii) faster learning about fundamentals. We combine these two channels to be able to preserve the optimal exit decisions of all entrants, allowing us to use observed life-cycle dynamics to evaluate it. In Section 5, we discuss the extent to which these two channels can be isolated.

#### 3.1.1 Accelerating organizational capital accumulation

The technology accelerates a firm’s organizational capital accumulation while young, without providing foresight about its mature fundamentals beyond what is revealed during youth in the factual economy. Specifically, it raises each firm’s productivity  $z_k$  while young to its expected average productivity while mature conditional on the information set available before reaching maturity (i.e.,  $\omega_\tau$ ). In other words, the path of productivities of every firm  $S$  while young becomes constant at  $\zeta_\tau(S)$ , as defined by Equation 2.4. Figure 2 illustrates simple examples of these counterfactual productivity paths.

We interpret this counterfactual as a reduced-form representation of a technological shift. While in practice firms may actively invest resources to build organizational capital, our model treats the productivity sequence as an intrinsic characteristic of the firm’s type. Consequently, we abstract away from the specific production function of learning

— e.g., whether the new technology acts as a substitute or complement to existing learning inputs, or how it alters the costs associated with knowledge accumulation. Instead, we focus solely on the output gains realized when a technology allows firms to bypass the initial period of potentially low productivity.

*AI and organizational capital.* An important constraint on organizational growth is the difficulty of transferring and preserving tacit knowledge — the “know-how” that resides in individuals and is difficult to articulate, store, or share (Polanyi, 1966; Nonaka, 1994; Garicano, 2000). This constraint may be particularly binding early in a firm’s life cycle, when organizational routines have not yet stabilized and learning often relies on noisy, ad hoc feedback rather than substantial accumulated experience and codified practices (e.g., Nelson and Winter, 1985; March, 1991; Grant, 1996). As a result, organizational capital traditionally accumulates slowly (Atkeson and Kehoe, 2005; Hsieh and Klenow, 2014).

AI has the potential to relax these constraints by allowing to convert dispersed tacit knowledge into scalable organizational assets. For example, Brynjolfsson et al. (2025) find that generative AI can capture the tacit behaviors of a firm’s top performers — such as nuanced communication styles and diagnostic intuition — which are typically difficult to codify in explicit rules. By disseminating these best practices via real-time suggestions, AI can effectively scale the knowledge of the firm’s most skilled agents across the organization.

Similarly, AI-powered enterprise knowledge platforms — such as *Glean* or *o9’s Digital Brain* — aggregate information from emails, documents, code repositories, dashboards, and operational logs, making past decisions, workflows, and problem-solving routines retrievable at low cost. This facilitates knowledge transmission, replication, and storage, helping young firms to scale effective practices before they are fully institutionalized. These platforms may also help slow down organizational forgetting (Argote et al., 1990) and reduce the erosion of managerial quality that typically occurs when practices are transmitted informally or left undocumented (Bloom et al., 2016). In addition, AI can accelerate early-stage feedback by allowing firms to extract more information from limited experience (Agrawal et al., 2022).

### 3.1.2 Accelerating learning about fundamentals

The technology also accelerates learning about fundamentals. It does so in two steps. First, it accelerates learning after entry: firms learn upon entry all the information encoded in  $\omega_\tau$  — which would have taken them  $\tau$  periods to learn in the factual economy. Second, it accelerates learning before entry. In particular, it allows potential entrants to

foresee which firms exit before reaching maturity in the factual economy. In other words, potential entrants can partition the set of all types  $\mathcal{S}$  into “bad types”  $\mathcal{S}_{\leq\tau}$  who exit before maturity and “good types”  $\mathcal{S}_{>\tau}$  that reach maturity in the factual economy.<sup>8</sup> In Section 6.2.2, we discuss the difficulty of considering these two steps of acceleration of learning about fundamentals in isolation.

*AI and learning about fundamentals.* We interpret this acceleration of learning about fundamentals as AI enabling entrants to direct their idea-generation effort toward searching within a particular subset of types. This form of acceleration captures AI’s growing ability to generate early, low-cost signals about a firm’s long-run prospects — signals that, absent AI, would only be revealed gradually through experience. In the factual economy, firms learn about their underlying type through realized performance: demand realizations, cost shocks, operational frictions, and survival itself serve as noisy and delayed signals about long-run fundamentals. As a result, both firms and investors must rely on prolonged experimentation to distinguish viable from non-viable business models.<sup>9</sup>

AI may relax this constraint by enabling potential entrants to simulate, forecast, and stress-test key aspects of their business before entry. A growing class of AI prediction and simulation tools — such as *Expected Parrot*, *Kruncher*, and *Shopify’s SimGym* — allow firms to obtain years of market feedback using synthetic environments and large-scale data. *Expected Parrot* uses agent-based consumer simulations to forecast demand and evaluate pricing or product designs; *Kruncher* automates due-diligence analytics in private markets to assess startup viability; and *SimGym* deploys AI-generated shoppers to simulate browsing and purchasing behavior, allowing firms to test product presentation and customer experience prior to interacting with real consumers.

These tools generate coarse but informative signals that help distinguish business models that are viable from those that are not. Consistent with this interpretation, recent evidence shows that AI agents already outperform experienced human venture capital analysts in predicting which early-stage startups survive ([Vismara et al., 2025](#)).

---

<sup>8</sup> The technology does not allow potential entrants to differentiate between different survivor types. This is consistent with the interpretation that it is trained on factual data about firm exits, but not underlying profitability. This assumption is conservative: If, in addition, the technology allowed potential entrants to know which among the types in  $\mathcal{S}_{>\tau}$  have the highest expected present value, then our VOLT estimates would be a lower bound on its value.

<sup>9</sup> Entrepreneurship training and accelerator programs are often explicitly designed not only to improve outcomes for high-potential ventures, but also to encourage “fast failure” among ventures with poor long-run prospects ([Cohen et al., 2019](#); [Bailey et al., 2026](#)).

### 3.2 VOLT Definition and Sufficient Statistic

We measure the Value of Organizational Learning Technologies (VOLT) as the increase in stationary equilibrium aggregate output that results from accelerating the first  $\tau$  years of learning. Formally,

$$\text{VOLT}_\tau := \frac{Y_\tau^c}{Y},$$

where  $Y$  is the stationary aggregate output in the factual economy and  $Y_\tau^c$  is the stationary aggregate output in the counterfactual economy with maturity threshold  $\tau$ .

For the result that follows, we define the statistics  $y_\tau$  and  $\ell_\tau$  as the ratio of mature to average firm size and lifespan in the factual economy. Specifically, the statistics are  $y_\tau := \bar{y}_{>\tau}/\bar{y}$  and  $\ell_\tau := \bar{\ell}_{>\tau}/\bar{\ell}$ , where  $\bar{y}$  and  $\bar{y}_{>\tau}$  are the average output of all firms and mature firms, respectively, while  $\bar{\ell}$  and  $\bar{\ell}_{>\tau}$  are their average lifespans.<sup>10</sup>

**Proposition 1. (VOLT Sufficient Statistic)** *The Value of Organizational Learning Technologies (VOLT) is determined by two statistics: the ratio of mature to average firm size,  $y_\tau$ , and the ratio of mature to average firm lifespan,  $\ell_\tau$ . In particular,*

$$\text{VOLT}_\tau = (y_\tau \cdot \ell_\tau)^{1-\nu} \tag{3.1}$$

given span of control parameter  $\nu \in (0, 1)$ .

A high value of  $y_\tau$  indicates that firms that survive beyond age  $\tau$  are substantially more productive than the average entrant. This gap can arise because organizational capital accumulates slowly — so firms spend many early years operating far below their mature productivity — or because survival is highly selective, so that only exceptionally productive firms reach maturity.

A high value of  $\ell_\tau$  reflects substantial early exit among young firms. This is symptomatic of an economy where entrants are not very effective at screening out bad types before entry. In such an environment, technologies that accelerate learning about fundamentals have large effects on aggregate output because they reallocate entrants toward firms with substantially longer lifespans.

*Proof of Proposition 1.* As is standard in models with decreasing returns to scale at the firm

<sup>10</sup> We refer to  $y_\tau$  as the ratio of mature to average firm *size* because the first-order condition 2.1 implies  $y_\tau = n_\tau = \bar{n}_{>\tau}/\bar{n}$ , where  $\bar{n}$  and  $\bar{n}_{>\tau}$  are the average employment of all firms and mature firms, respectively. This equivalence may not hold when marginal products of labor are not equalized across firms, a question we return to in Section 5.2.

level, we can aggregate the firm-level production function (2.1) to obtain:<sup>11</sup>

$$Y = (ZM)^{1-\nu} N^\nu, \quad (3.2)$$

where  $M$  is the mass of active firms in stationary equilibrium, and  $Z \equiv \frac{1}{M} \int_0^M z_i^{\frac{1}{1-\nu}} di$  is the cross-sectional average of firm productivity  $z^{\frac{1}{1-\nu}}$ . The product  $ZM$  can be understood as the aggregate stock of organizational capital in stationary equilibrium: it is the average organizational capital per firm,  $Z$ , times then number of active firms,  $M$ .

Given that employment  $N$  is fixed (see Section 2.4), aggregate output comparisons reduce to comparisons of organizational capital. In particular:

$$\text{VOLT}_\tau \equiv \frac{Y_\tau^c}{Y} = \left( \frac{Z_\tau^c}{Z} \cdot \frac{M_\tau^c}{M} \right)^{1-\nu}.$$

Computing  $\text{VOLT}_\tau$  therefore requires evaluating only two objects: the ratio of average productivities  $\frac{Z_\tau^c}{Z}$  and the ratio of firm masses  $\frac{M_\tau^c}{M}$  in stationary equilibrium.

The key observation that leads to Proposition 1 is that firm exit choices *of the firms that enter the counterfactual* remain exactly as in the factual. This allows us to express the ratio of counterfactual to factual aggregates using the factual ratios  $y_\tau$  and  $\ell_\tau$ .

**Step 1: Invariance of Exit Choices.** We start by considering the direct effect of the technology *holding the price  $p$  fixed*. First, note that only survivor types enter in the counterfactual. The technology allows potential entrants to distinguish between the set of exitors  $\mathcal{S}_{\leq \tau}$  and survivors  $\mathcal{S}_{> \tau}$ . Given Assumption 3, potential entrants find it optimal to select the “good” set  $\mathcal{S}_{> \tau}$ , and draw types from it alone in the counterfactual. This optimality relies on both parts of Assumption 3: (i) the higher mature productivity guarantees that, in the counterfactual, survivors generate higher flow profits than exitors while young, and (ii) the higher mature value ensures that the survivors’ expected lifespans are long enough to generate more lifetime value than exitors in the counterfactual.

Second, note that exit decisions while mature in the counterfactual are exactly the same as in the factual. By definition, the technology does not change the states  $S_k$  beyond age  $\tau$  for any type  $S$ . The technology does reveal upon entry the expected average productivity  $\zeta_\tau(S)$  that a firm has after age  $\tau$  in the factual. However, by Assumption 1, this information does not add to the information encoded in states  $S_k$  of mature firms with age  $k > \tau$  because, in the factual, the firm could already compute  $\zeta_\tau(S)$  at time  $\tau$ .

---

<sup>11</sup> Aggregating the first order condition 2.1 gives  $N = (p\nu)^{\frac{1}{1-\nu}} ZM$  and  $Y = (p\nu)^{\frac{\nu}{1-\nu}} ZM$ . Combining these two conditions gives Equation 3.2.

Third, note that survivors still do not exit while young in the counterfactual. By definition, survivors (the only types that enter in the counterfactual) never exit while young in the factual economy. In particular, at age  $\tau$ , these types decide to stay in the market, so they have a positive expected present value  $\mathbb{E}[V(S_{\tau+1}) | S_\tau]$ . In the counterfactual, their average output at every age  $k \leq \tau$  is equal to their expected average output after  $\tau$ . In particular, the profits net of fixed costs  $(\bar{\pi}_{>\tau}(S_\tau) - f)$  are weakly positive at every age  $k \leq \tau$ . Moreover, as just explained above, their expected present value at age  $\tau$  in the counterfactual is the same as in the factual. Given that firms learn the information in  $\omega_\tau$  upon entry in the counterfactual, they know the state  $S_\tau$  at entry as well, so the expected present value at any age  $k \leq \tau$  is

$$V^c(S_k) = \sum_{a=k}^{\tau} \beta^{a-k} \underbrace{(\bar{\pi}_{>\tau}(S_\tau) - f)}_{\geq 0} + \beta^{\tau-k} \underbrace{\mathbb{E}[V(S_{\tau+1}) | S_\tau]}_{\geq 0} \geq 0.$$

Hence, in the counterfactual, firms that enter do not exit while young.

To conclude this first step of the proof, it only remains to consider the effects of the equilibrium change in the price  $p$ . Accelerating learning raises aggregate productivity, increasing output and thus, per Equation 2.5, lowering the equilibrium output price  $p$ . We might worry that, by reducing variable profits, lower prices could force some firms in the counterfactual to exit earlier than they do in the factual, thereby violating invariance of exit choices. However, Assumption 2 guarantees that the fixed costs fall in the same proportion ( $p^{\frac{1}{1-\nu}}$ ) as the variable profits of mature firms. As a result, value functions are also proportional to  $p^{\frac{1}{1-\nu}}$ , so their sign is preserved and exit choices remain as in the factual.

**Step 2. Mapping Counterfactual Aggregates to Factual Data.** The two variables that determine aggregate output are the mass of firms  $M$  and the aggregate productivity  $Z$ . Regarding the mass of firms, given the fixed unit flow of entrants, the mass of active firms  $M$  equals the expected lifespan both in the factual and counterfactual. As we have just shown above, only survivor types are present in the counterfactual. Therefore, the ratio  $M_\tau^c / M$  is equal to  $\bar{\ell}_{>\tau} / \bar{\ell}$ .

Regarding aggregate productivity, in a stationary equilibrium, the cross-sectional average  $Z \equiv \frac{1}{M} \int_0^M z_i^{\frac{1}{1-\nu}} di$  is equal to the longitudinal expectation  $\mathbb{E}_\mu \left[ \frac{1}{\ell(S)} \sum_{k=0}^{\ell(S)-1} z_k(S)^{\frac{1}{1-\nu}} \right]$ . Thus, by construction,  $Z_\tau^c = \mathbb{E}_\mu [\zeta_\tau(S) | S \in S_{>\tau}]$ . Combining these equations with the first-order condition 2.1 and the fact that the cross-sectional distribution of firms at any

point in time mirrors the dynamic evolution of a single cohort, we obtain

$$Z = \mathbb{E}_\mu \left[ \frac{1}{\ell(S)} \sum_{k=0}^{\ell(S)-1} z_k(S)^{\frac{1}{1-\nu}} \right] = (\nu p)^{\frac{\nu}{1-\nu}} \bar{y} \text{ and } Z_\tau^c = \mathbb{E}_\mu [\zeta_\tau(S) \mid S \in S_{>\tau}] = (\nu p)^{\frac{\nu}{1-\nu}} \bar{y}_{>\tau}$$

We conclude that  $Z_\tau^c/Z = y_\tau$ . □

The sufficient statistics  $y_\tau$  and  $\ell_\tau$  capture the two channels through which organizational learning technologies increase aggregate output: by raising average per-period productivity and by extending firms' effective lifespans. Both channels operate through an increase in aggregate organizational capital. The terms  $y_\tau$  and  $\ell_\tau$  therefore enter Equation 3.1 with exponent  $1 - \nu$ , reflecting decreasing returns to organizational capital (see Equation 3.2)

The proof of Proposition 1 makes transparent the role of Assumptions 1–3. Assumption 1 guarantees that the counterfactual disclosure of  $\zeta_\tau(S)$  upon entry does not change what a firm would already infer by time  $\tau$ . Without this, young survivor types could potentially make different exit decisions in the factual and counterfactual. Assumption 2 ensures that exit choices of mature firms do not depend on the particularities of how young firms build organizational capital. Without it, the exit behavior of mature firms could differ in the factual and counterfactual economies even if their productivity paths are the same in both economies.

Finally, Assumption 3 ensures that, once the technology separates exitors from survivors, entrants strictly prefer to draw from the survivor set. Without this assumption, entrants could instead select exitor types, and the aggregate implications would hinge on counterfactual objects we do not observe —  $\zeta_\tau(S)$  and mature continuation/exit behavior for types that exit before reaching maturity in the factual. In Section 5, we describe how Proposition 1 changes in different natural extensions of our baseline setting.

## 4 Measuring VOLT

We now detail the data and methodology we use to estimate VOLT (Section 4.1), and provide our baseline VOLT estimates (Section 4.2).

### 4.1 Data and Measurement

We estimate VOLT using the Business Dynamics Statistics (BDS) of the U.S. Census Bureau. This is a publicly available dataset covering the entire U.S. private sector since 1978.

Constructed from the confidential micro-data of the Longitudinal Business Database, the BDS aggregates firm and establishment-level records into annual series on counts, employment, and exits. Following the standard approach in the literature (e.g., [Atkeson and Kehoe, 2005](#)), we focus on establishment — rather than firm — data. This is consistent with the interpretation that organizational learning is plant specific. Our baseline estimates are based on the most recent data available, corresponding to 2023.

A limitation of the BDS is that it does not contain direct measures of output. Accordingly, we proxy the size ratio  $y_\tau$  with the employment ratio  $n_\tau$ . This mapping is natural in our baseline framework because  $y_\tau = n_\tau$ . This equivalence may fail in the presence of wedges that prevent the equalization of marginal products of labor across establishments, an issue we discuss in [Section 5.2](#).

The structure of the BDS data determines the granularity of our analysis. At the economy-wide level, establishments are grouped into twelve age bins: single-year bins for ages 0 through 5; five-year bins for the years 6 to 25; a “26 years and over” category; and a “Left Censored” bin for establishments born before 1977. In 2023, “born before 1977” corresponds to ages above 46. Hence, this granularity allows us to identify mature establishments for the maturity thresholds 0, 1, 2, 3, 4, 5, 10, 15, 20, 25 and 46. At the 3-digit NAICS level, the data are broken down by establishment age into five bins: “less than one year old”, “1–5 years,” “6–10 years,” “11+ years,” and “unknown age” (born before 1977). This granularity allows us to identify mature establishments for mature thresholds 0, 5, 10 and 46.

The BDS reports, for each bin and year, the stock of establishments active, their total employment, and the flow of establishments that exit in that year. We estimate the average employment  $\bar{n}$  and the average mature employment  $\bar{n}_{>\tau}$  simply by dividing the total employment over the number of establishments, and the total employment in mature establishments over the number of mature establishments, respectively.

Estimating the average firm lifespan  $\bar{\ell}$  and average mature lifespan  $\bar{\ell}_{>\tau}$  is more involved because we don’t observe them directly. Instead, we have to infer them from exit data. We use the following synthetic cohort approach. Let  $\mathcal{B} = \{b_1, b_2, \dots, b_K\}$  denote the set of age bins provided in the data, where the final bin  $b_K$  is open-ended (e.g., establishments born before 1977). For each bin  $b_j$ , let  $\delta_j$  denote the annual exit hazard, calculated as the ratio of exits to active establishments in that bin:

$$\delta_j = \frac{\text{Exits}_j}{\text{Active}_j}.$$

We assume a constant hazard rate within bins. In particular, we define the age-specific

hazard rate  $\delta(k)$  simply as the hazard of the bin that age  $k$  falls into. That is,  $\delta(k) = \delta_j$  if age  $k$  is within the range of bin  $b_j$ .<sup>12</sup>

Let  $t_K$  be the starting age of the last, open ended, bin (using 2023 data, this is  $t_K = 47$ ). For each  $t \leq t_K$ , we construct the survival function  $P(t)$  — representing the probability that an establishment survives to age  $t$  — recursively:

$$P(0) = 1 \text{ and } P(t) = \prod_{k=0}^{t-1} (1 - \delta(k)), \text{ for all } 1 \leq t \leq t_K$$

The aggregate lifespan  $\bar{\ell}$  is the sum of survival probabilities,  $\sum_{t=0}^{\infty} P(t)$ .<sup>13</sup> We estimate it as:

$$\sum_{t=0}^{t_K-1} P(t) + \frac{P(t_K)}{\delta_K},$$

where the contribution of the open-ended tail starting at age  $t_K$  with constant hazard  $\delta_K$  is the sum of an infinite geometric series:

$$\sum_{k=0}^{\infty} P(t_K + k) = P(t_K) \sum_{k=0}^{\infty} (1 - \delta_K)^k = \frac{P(t_K)}{\delta_K}.$$

Similarly, the expected lifespan of a mature establishment,  $\bar{\ell}_{>\tau}$ , is the youth lifespan ( $\tau + 1$ ) plus the sum of (conditional) survival probabilities beyond youth:

$$\tau + 1 + \frac{1}{P(\tau + 1)} \left( \sum_{t=\tau+1}^{t_K-1} P(t) + \frac{P(t_K)}{\delta_K} \right)$$

In particular, for the maximum mature threshold allowed by the BDS data (i.e.,  $t_K = \tau + 1$ ), this is  $\tau + 1 + \frac{1}{\delta_K}$ . Section B of the Appendix reports our estimates of  $\bar{n}_{>\tau}$ ,  $n_\tau$ ,  $\bar{\ell}_{>\tau}$  and  $\ell_\tau$  for all maturity thresholds allowed by the BDS data.

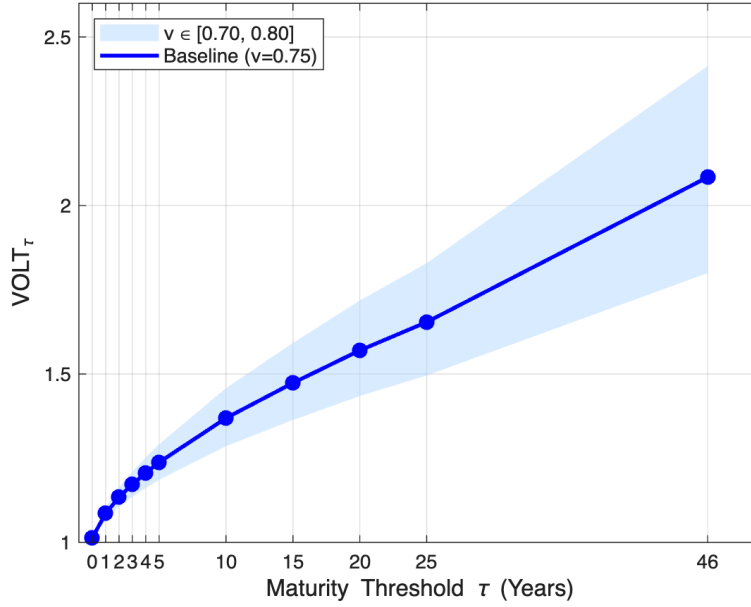
## 4.2 VOLT Estimates

Figure 3 illustrates the estimated  $\text{VOLT}_\tau$  in the U.S. for all values of the maturity threshold  $\tau$  supported by the BDS.<sup>14</sup> These estimates reveal that the aggregate gains from organizational learning are substantial even for modest learning accelerations. Accelerating just the first 2 years of learning ( $\tau = 1$ ) increases steady-state output by approximately 9%.

<sup>12</sup> The BDS does not report the number of exits of firms less than one year old; we assume the hazard rate of that bin is 0.

<sup>13</sup> In line with our model, this assumes that firms that exit at age  $k$  have a lifespan of  $k + 1$ .

<sup>14</sup> Section B of the Appendix reports the corresponding estimates.




**Figure 3:** VOLT as a function of the maturity threshold  $\tau$ . The solid line reports the baseline estimates assuming span of control parameter  $\nu = 0.75$ . The shaded region indicates the range of estimates for  $\nu \in [0.70, 0.80]$ .

The gains accumulate rapidly in the early years: accelerating the first six years ( $\tau = 5$ ) raises output by nearly 24%, and accelerating the first 11 years ( $\tau = 10$ ) yields an increase of roughly 37%. When the maturity threshold is the maximum supported by the BDS data ( $\tau = 46$ ), the value of the technology is of the order of one GDP ( $\text{VOLT}_{46} \approx 2.08$ ), suggesting that accelerating learning has the potential to double aggregate output. These magnitudes remain economically significant across a reasonable range of span-of-control parameters.

Figure 4 decomposes the sources of  $\text{VOLT}_{\tau}$ , revealing that the lifespan effect  $\ell_{\tau}^{1-\nu}$  acts as its primary driver. When the mature threshold is  $\tau = 46$ , mature establishments have an expected lifespan about seven times that of the average firm ( $\ell_{46} \approx 7$ ), and are about 3 times as productive as the average establishment ( $y_{46} \approx 2.7$ ). As a result, the increase in aggregate output driven by the lifespan effect  $\ell_{46}^{1-\nu}$  alone is 62% of the total increase suggested by  $\text{VOLT}_{46}$ , while the size effect  $y_{46}^{1-\nu}$  alone is 28% of the total increase suggested by  $\text{VOLT}_{46}$ . The remaining share of the output increase suggested by  $\text{VOLT}_{46}$  is due to the interaction between the two effects.<sup>15</sup>

These estimates highlight the quantitative importance of the lifespan effect  $\ell_{\tau}^{1-\nu}$ . Standard analyses of aggregate productivity typically focus on the average efficiency of active

<sup>15</sup> Section B of the Appendix reports these shares for all maturity thresholds  $\tau$ .



**Figure 4:** Decomposition of VOLT into its components  $y^{1-\nu}$  and  $l^{1-\nu}$  as a function of the maturity thresholds  $\tau$  allowed by the BDS data. In this figure,  $\nu = 0.75$ .

firms,  $Z$ , which in our interpretation corresponds to average organizational capital *per firm*. However, as emphasized by [Hopenhayn \(2014\)](#), aggregate output depends not only on average productivity per firm but also on the total mass of active firms,  $M$ . In particular, an economy may display a low Solow residual —  $Y/N$  in our setting — either because firms have accumulated little organizational capital or because few firms are active. We return to this distinction in Sections 5.1 and 6.3, where we explore alternative assumptions about entry and discuss the relationship between the lifespan effect and business dynamism (e.g., [Aghion and Howitt, 1990](#)).

We interpret the dominance of the lifespan statistic  $l_\tau$  with caution. At first glance, the fact that  $l_\tau$  exceeds  $y_\tau$  could be read as evidence that selection is the primary source of value creation. This interpretation is misleading. The lifespan statistic is an equilibrium outcome shaped by both learning channels. To see this, note that if the learning technology only accelerated the accumulation of organizational capital, it could endogenously reduce exit rates, thus increasing longevity and generating a lifespan effect. Thus, the finding that  $l_\tau > y_\tau$  does not imply that learning about fundamentals is the dominant mechanism; instead, it indicates that, in the class of models we consider, a central manifestation of accelerated learning — regardless of the specific learning channel — is greater firm longevity. We return to this distinction in Section 6.2.1, where we discuss the extent to which we can disentangle the contributions of each learning channel.

### 4.3 Industry VOLT Estimates

Our aggregate VOLT estimates potentially mask substantial heterogeneity across industries. To quantify this heterogeneity, we take our sufficient-statistic result (Equation 3.1) to data disaggregated at the 3-digit NAICS level. We treat each industry  $j$  as a distinct production environment governed by the firm dynamics described in Section 2, and we assume the final good is produced by a Cobb-Douglas aggregator of industry outputs. This implies a unitary elasticity of substitution across industries, ensuring that the expenditure share — and thus the labor share — of each industry remains constant in equilibrium. In this case, the value of organizational learning technologies for each industry  $j$  is given by:

$$\text{VOLT}_{\tau,j} := \frac{Y_{\tau,j}}{Y_j} = (y_{\tau,j} \ell_{\tau,j})^{1-\nu} \text{ where } y_{\tau,j} := \frac{\bar{y}_{>\tau,j}}{\bar{y}_j} \text{ and } \ell_{\tau,j} := \frac{\bar{\ell}_{>\tau,j}}{\bar{\ell}_j}$$


The industry statistics  $y_{\tau,j}$  and  $\ell_{\tau,j}$  have interpretations analogous to their aggregate counterparts. Figure 5 illustrates our VOLT estimates for all NAICS-3 industries in the U.S.<sup>16</sup> The value of organizational learning technologies exhibits significant heterogeneity (Avg.  $\text{VOLT}_{46} \approx 1.96$ ,  $\sigma \approx 0.39$ ). Estimates range from a low of 1.27 to a high of 3.41, suggesting that while the value of organizational learning technologies is pervasive, its magnitude depends heavily on the specific industry.

Decomposing these estimates reveals that this variation is driven by heterogeneity in both organizational lifespan,  $\ell_{\tau}$ , and firm size,  $y_{\tau}$ . High-VOLT industries generally exhibit a combination of longevity and large scale, though the primary driver varies by sector. The potential gains are highest in Couriers & Messengers ( $\text{VOLT}_{46} \approx 3.4$ ), which benefits from high values in both lifespan ( $\ell_{46} \approx 14$ ) and scale ( $y_{46} \approx 9.6$ ). In contrast, the high value in Apparel Manufacturing ( $\text{VOLT}_{46} \approx 2.9$ ) is driven disproportionately by its exceptional organizational lifespan ( $\ell_{46} \approx 22$ ), whereas Performing Arts ( $\text{VOLT}_{46} \approx 2.9$ ) relies on a balanced contribution from both longevity and size.<sup>17</sup>

We now aggregate industry-level VOLT estimates assuming a Cobb–Douglas production function for the final good. The resulting aggregate VOLT is the geometric average

<sup>16</sup> Section C of the Appendix reports the corresponding estimates. We exclude four sectors from the analysis — Fishing, Hunting and Trapping (NAICS 114), Pipeline Transportation (NAICS 486), Monetary Authorities (NAICS 521) and Funds, Trusts and other Financial Vehicles (NAICS 525) — because insufficient data on mature exits precludes the calculation of finite lifespan statistics.

<sup>17</sup> These industry-level VOLT estimates are naturally sensitive to the elasticity of demand. We return to this issue in Section 5.3.



**Figure 5:**  $\log(VOLT_{46})$  for each NAICS 3-digit sector decomposed into the size effect (dark blue,  $(1 - \nu) \ln(y_{46})$ ) and lifespan effect (light blue,  $(1 - \nu) \ln(\ell_{46})$ ). In this figure,  $\nu = 0.75$ .

of industry values:

$$\text{VOLT}_{\tau,\text{agg}} := \prod_{j=1}^N (\text{VOLT}_{\tau,j})^{\theta_j}$$

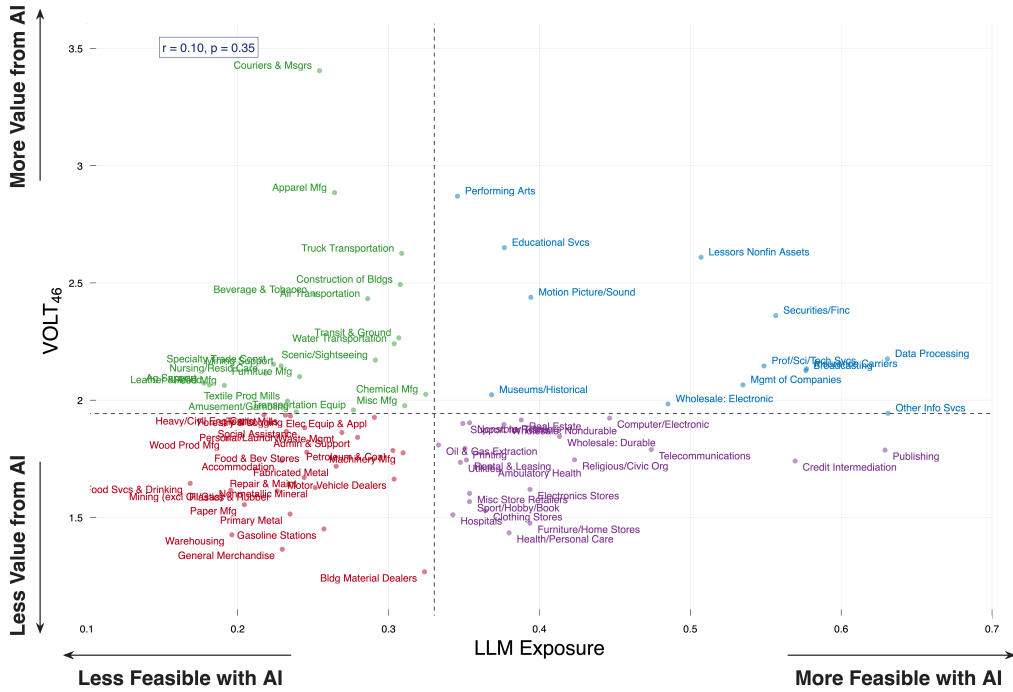
We estimate industry shares  $\theta_j$  using nominal value-added shares from the Bureau of Economic Analysis (BEA) GDP by Industry accounts for 2023, matching BEA industry categories to their corresponding NAICS industries in the BDS.<sup>18</sup> This yields an economy-wide VOLT of 1.78. The modest difference with the aggregate 2.08 estimate is expected: because the gains from organizational learning are concave (governed by  $1 - \nu$ ), a single-industry model that calculates gains based on aggregate statistics mechanically yields a higher estimate than an aggregation of industry gains. That the two estimates are close suggests that our central finding — that organizational learning technologies have the potential to substantially increase aggregate output — is not driven by a small number of outlier industries.

#### 4.4 VOLT vs LLM Exposure

A growing literature is developing measures of AI exposure to assess where the capabilities of artificial intelligence overlap with the tasks performed in different occupations and industries (Brynjolfsson et al., 2018; Eloundou et al., 2024; Labaschin et al., 2025). In contrast to VOLT, these exposure metrics are designed to capture the *technical feasibility* of adopting AI rather than measuring the potential value of adopting AI. In this section, we show that they are largely uncorrelated to VOLT.

We focus on the occupation-level large-language model (LLM) exposure indices developed by Eloundou et al. (2024). These indices are constructed using detailed task descriptions from the O\*NET database, which are independently evaluated by both human annotators and GPT-4. The evaluation rubric assesses whether access to a model with GPT-4-level capabilities would reduce the time required to complete a task by at least 50 percent while maintaining equivalent quality. Tasks are categorized into two groups: *directly exposed tasks*  $E_1$ , which can be substantially accelerated by an LLM operating on its own, and *indirectly exposed tasks*  $E_2$ , for which similar gains are achievable only when the LLM is combined with task-specific complementary software. Following the baseline convention in this literature, we define occupation  $o$ 's exposure as  $E_o = \mathbb{E} [E_1 + 0.5 \times E_2]$ , where the expectation is taken over tasks within the occupation.

<sup>18</sup> We resolve the many-to-many BEA–NAICS correspondence by working on the coarsest common industry partition (aggregating NAICS industries when BEA is coarser, and summing BEA lines when BEA is finer), restricting to the set of industries that can be consistently matched to the BDS, and renormalizing shares to sum to one.



**Figure 6:** VOLT<sub>46</sub> vs. LLM Exposure (Eloundou et al., 2024) by 3-Digit NAICS sector. Section C in the Appendix reports the underlying values.

To characterize exposure across industries, we aggregate occupation-level exposure using employment shares from the Bureau of Labor Statistics’ Occupational Employment and Wage Statistics (OES). For each industry  $s$ , we compute an employment-weighted average of the exposure of occupations employed in that industry:

$$E_s = \sum_o \left( \frac{\text{Emp}_{o,s}}{\text{Emp}_s} \right) E_o$$

where  $\text{Emp}_{o,s}$  denotes employment of occupation  $o$  in industry  $s$  as measured in the OES, and  $\text{Emp}_s$  is total employment in industry  $s$ . We emphasize that the LLM Exposure Index measures *technical potential* rather than realized outcomes.

Figure 6 plots VOLT<sub>46</sub> against LLM exposure by sector. The correlation between the two measures is weak (0.1) and statistically insignificant (the p-value is 0.35). The weak correlation suggests that VOLT and exposure pick up distinct dimensions of AI’s impact. Sectors in the top-right quadrant represent the frontier of potential AI transformation, combining high VOLT with high exposure to LLM automation. This group includes Educational Services (VOLT<sub>46</sub> ≈ 2.65, LLM Exposure ≈ 0.38) and Securities & Financial Investments (VOLT<sub>46</sub> ≈ 2.36, LLM Exposure ≈ 0.56). Similarly, Data Processing and Hosting

has strong value of organizational learning technologies ( $VOLT_{46} \approx 2.18$ ) and high LLM exposure (0.63).

In contrast, the top-left quadrant highlights industries where the value of accelerating learning is substantial, yet current LLM capabilities offer limited ability to automate the relevant tasks. This disconnect is most visible in Couriers and Messengers, which has the highest VOLT in the dataset ( $VOLT_{46} \approx 3.41$ ) driven by massive size and lifespan effects ( $y_{46} \approx 9.6$ ,  $\ell_{46} \approx 14$ ), yet has low LLM exposure (0.25) probably due to the physical nature of logistics. Apparel Manufacturing fits a similar profile; its exceptional organizational lifespan statistic ( $\ell_{46} \approx 22$ ) drives a high VOLT of 2.88, but presumably its physical core tasks keep LLM exposure low (0.26), suggesting these sectors require robotics rather than language models to capture their potential efficiency gains.

## 5 Extensions

This section extends the baseline framework by considering alternative entry assumptions (Section 5.1), age-dependent labor wedges (Section 5.2), and general demand elasticity (Section 5.3).

### 5.1 Entry

In the baseline class of models of Section 2, the flow of entrants is exogenous, so faster learning does not change the equilibrium entry rate. We now show that the learning technology raises the expected value of entry. This implies that if the flow of entrants reacted to the value of entry — as in [Hopenhayn \(1992\)](#), [Hopenhayn and Rogerson \(1993\)](#) or [Melitz \(2003\)](#) — our baseline VOLT estimates would be biased downward. We also consider an alternative environment where — as in [Atkeson and Kehoe \(2005\)](#) — the mass of firms is fixed; in this case, the aggregate implications of the learning technology are purely driven by total factor productivity (TFP).

#### 5.1.1 The Learning Technology Increases the Value of Entry

The learning technology has opposing effects on the net present value of entry. On the one hand, it increases the value of entry by raising productivity while young and reducing the fixed cost of operation. On the other hand, it lowers the value of entry by reducing the output price. We now show that the former effect dominates. The central part of the argument is showing that total (undiscounted) lifetime profits of a cohort increase.

Total undiscounted lifetime profits are equal to the aggregate net profits (variable profits minus fixed costs) generated by the cross-section of active firms in stationary equilibrium. Aggregate variable profits are constant because, in this class of models, aggregate expenditure is fixed at  $E$ , and variable profits are a constant share  $\nu$  of expenditure.

Aggregate fixed costs equal the mass of active firms times the fixed cost per firm,  $Mf$ . The learning technology increases lifespans, so the mass of active firms rises by a factor  $\ell_\tau$ . The technology also raises aggregate output (by a factor of  $(y_\tau \ell_\tau)^{1-\nu}$ , as shown by Proposition 1), which forces the equilibrium price to decline in order to keep total expenditure at  $E$ . Given that the per-firm fixed cost  $f$  is indexed to mature-firm profits, it falls in proportion to the equilibrium price level; in particular it declines by a factor  $1/(y_\tau \ell_\tau)$ . Multiplying the increased mass of firms by the reduced fixed cost per firm implies that total fixed costs scale by

$$M^c f^c \propto \ell_\tau \cdot \frac{1}{y_\tau \ell_\tau} = \frac{1}{y_\tau}$$

Hence, under the natural assumption that mature firms are at least as productive as the average firm ( $y_\tau \geq 1$ ), total fixed costs decrease. Because total variable profits are unchanged while total fixed costs fall, stationary aggregate net profits increase, and hence so does each cohort's undiscounted lifetime profit flow.

As we show in Appendix A, under the natural assumption that (i) average productivity among young firms is weakly increasing with age, and (ii) the expected average productivity of a firm after it matures,  $\zeta_\tau$ , is greater than its average productivity while young, this implies that the net present value of entry goes up for any discount factor  $\beta \leq 1$ . If entry were elastic rather than fixed, this increase in the net present value of entry would translate into a higher entry flow, yielding an even larger mass of active firms in the counterfactual economy. Consequently, VOLT in this extension would be larger than we have estimated under the baseline class of models in Section 2.

### 5.1.2 Alternative Market Structure: Fixed Number of Firms

Rather than fixing the flow of entrants, we can consider the polar opposite case in which the number of active firms is fixed. For example, in Atkeson and Kehoe (2005), the economy is endowed with a fixed stock of managers, each able to operate one firm. Hence, the total mass  $M$  of firms is fixed, and the managers' wage (the fixed cost of operation) adjusts endogenously to drive the expected present value of entry to zero.

In this alternative environment, because the total mass of firms is constrained by the fixed supply of managers, the lifespan effect  $\ell_\tau$  — which relies on an expansion in the

number of active firms — is neutralized. The value of the learning technology is therefore solely driven by the increase in total factor productivity ( $Z$ ). However, this productivity effect can be amplified by the least productive firms exiting earlier than in the factual economy due to the following two reasons.

First, because the expected value of operating a firm rises, competition for the fixed number of managers drives up the equilibrium managerial compensation. Second, the increase in aggregate productivity drives down the equilibrium price. The combination of higher managerial wages and lower prices may force the least productive firms to exit earlier in the counterfactual than in the factual. This truncates the lower tail of the productivity distribution, thereby increasing TFP beyond the size effect  $y_\tau^{1-\nu}$ . Hence, in this alternative environment with a fixed number of firms, the value of organizational learning is bounded below by the TFP gains — captured by the size effect  $y_\tau^{1-\nu}$  — obtained in the class of models in Section 2.

## 5.2 Wedges

As evidenced by the first-order condition 2.1, employment and output in our baseline model are a direct reflection of productivity. This relationship allows us to identify the TFP effect of organizational learning technologies ( $Z^c / Z$ ) from the ratio of mature to average output  $y_\tau$ , or, equivalently, the ratio of mature to average employment,  $n_\tau$ . However, this mapping between output, employment and productivity is sensitive to wedges in firms' marginal product of labor that act as a firm-specific tax or subsidy on labor input. If such wedges are systematically correlated with firm age, our baseline estimates of the TFP effect of organizational learning technologies may be biased.

Age-dependent wedges may arise in a wide range of environments (Bergquist et al., 2026). For example, they can originate from financial frictions that constrain new or unproven firms but relax as firms build a track record (e.g., Clementi and Hopenhayn, 2006; Buera et al., 2011), size-dependent regulations that disproportionately burden young firms (e.g., Clementi and Hopenhayn, 2006; Guner et al., 2008; Levy, 2010; Buera et al., 2011), and adjustment costs that are more salient early in a firm's life cycle (e.g., Hopenhayn and Rogerson, 1993; Hopenhayn, 2014). Age-dependent wedges can also arise from standard forms of learning by doing (e.g., Arrow, 1962), which can be interpreted as an implicit labor subsidy that fades as the firm matures, or from distortions that may intensify with age, such as increasing market power (e.g., De Loecker et al., 2020; Akcigit et al., 2023).

To investigate how our VOLT estimates change in the presence of such wedges, we extend the class of models in Section 2 to allow for labor distortions that may increase

or decrease over the life cycle. In particular, we introduce age-specific labor wedges  $\eta_k$ , so the perceived wage of firms of age  $k$  is  $\eta_k$  times the market wage (which we have normalized to 1). Hence,  $\eta_k > 1$  corresponds to an age-specific tax on labor input, while  $\eta_k < 1$  corresponds to an age-specific subsidy. Firms optimal labor demand and output are given by the following modification of the firms' first order condition 2.1:

$$n^*(z_k) = \left( \frac{pv}{\eta_k} z_i \right)^{\frac{1}{1-\nu}} \quad \text{and} \quad y^*(z_k) = \frac{\eta_k}{pv} n^*(z_k). \quad (5.1)$$

In this distorted economy, the observed output and employment ratios,  $y_\tau$  and  $n_\tau$ , no longer coincide, and both conflate true productivity differences with wedges. Given that we don't observe firms' output in the BDS data, in the following discussion we focus on the employment ratio  $n_\tau$ . Also, to clarify in the simplest possible way how the statistic  $n_\tau$  may conflate productivity with wedges, we focus on the case in which  $\eta_k = \eta_{\leq \tau}$  for young firms ( $k \leq \tau$ ) and  $\eta_k = \eta_{> \tau}$  for mature firms ( $k > \tau$ ).

We define *Adjusted VOLT* $_\tau$  as the increase in aggregate stationary output  $Y/Y_\tau^c$  created by the accelerated learning technology described in Section 3 *while keeping age-dependent wedges fixed*. Letting  $s_{\leq \tau}$  and  $s_{> \tau}$  be the factual employment shares of young and mature firms, respectively, the relationship between the measured employment gap  $n_\tau$  and the productivity gap is:<sup>19</sup>

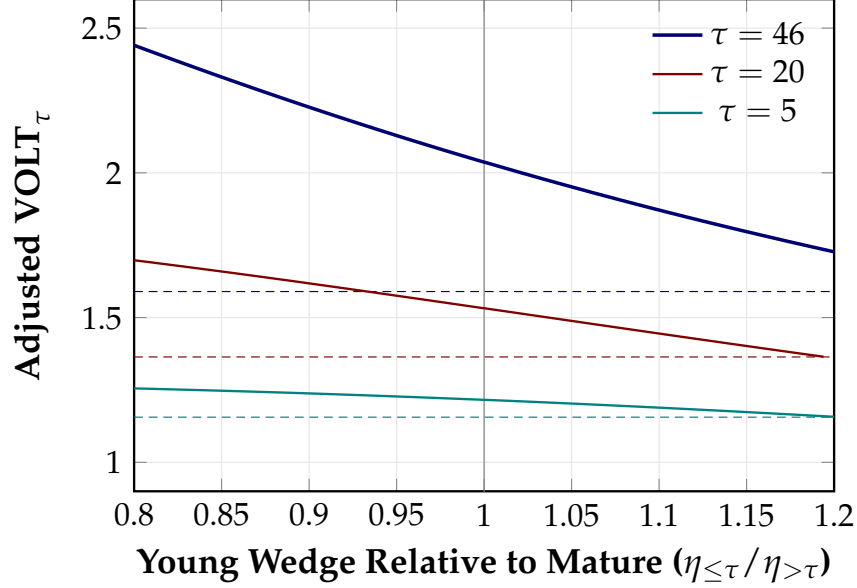
$$n_\tau = \frac{\bar{n}_{> \tau}}{\bar{n}} = \left( \frac{\bar{z}_{> \tau} \bar{\eta}_\tau}{\bar{z} \eta_{> \tau}} \right)^{\frac{1}{1-\nu}} \quad \text{where} \quad \bar{\eta}_\tau \equiv \left[ s_{\leq \tau} \cdot \eta_{\leq \tau}^{\frac{1}{1-\nu}} + s_{> \tau} \cdot \eta_{> \tau}^{\frac{1}{1-\nu}} \right]^{1-\nu} \quad (5.2)$$

Hence, the value of organizational learning technologies in the presence of age-dependent wedges is equal to our baseline VOLT estimate times the ratio of mature to average wedges:

$$\text{Adjusted VOLT}_\tau = \frac{\eta_{> \tau}}{\bar{\eta}_\tau} \text{VOLT}_\tau$$

In particular, if young firms face larger wedges than mature firms, then  $\text{VOLT}_\tau$  overestimates the value of organizational learning technologies. However, under the natural assumption that the average productivity of mature firms is at least as large as that of young firms,  $\bar{z}_{> \tau} \geq \bar{z}$ , it follows from Equation 5.2 that wedges can at most neutralize the size effect,  $\eta_{> \tau} / \bar{\eta}_\tau \geq 1 / n_\tau^{1-\nu}$ , so *Adjusted VOLT* $_\tau$  is at least as large as the lifespan effect  $\ell_\tau^{1-\nu}$ . Figure 7 illustrates, for several values of  $\tau$ , *Adjusted VOLT* $_\tau$  and its lower bound as a function of relative distortions  $\eta_{\leq \tau} / \eta_{> \tau}$ .

<sup>19</sup> From the static first order condition,  $n(z_k, \eta_k) = (pvz_k / \eta_k)^{\frac{1}{1-\nu}}$ , so group-average employments satisfy  $\bar{n}_{> \tau} \propto (\bar{z}_{> \tau} / \eta_{> \tau})^{-\frac{1}{1-\nu}}$  and  $\bar{n} \propto (\bar{z} / \bar{\eta}_\tau)^{-\frac{1}{1-\nu}}$ .



**Figure 7:** Adjusted  $\text{VOLT}_\tau$  as a function of the relative distortion facing young firms  $\eta_{\leq\tau}/\eta_{>\tau}$  for different values of the mature threshold  $\tau$ . The horizontal dashed lines indicate the lifespan effect,  $\ell_\tau^{1-\nu}$ , which is a lower bound on Adjusted VOLT provided that mature firms are, on average, at least as productive as the population average ( $\bar{z}_{>\tau} \geq \bar{z}$ ).

### 5.3 Demand elasticity

In the baseline class of models described in Section 2, aggregate expenditure is fixed at  $E$ . As we discussed in Section 2.4, this admits two natural interpretations: an aggregate economy with a fixed labor supply or an industry or small open economy within a larger economy facing a unit-elastic demand curve (e.g., a Cobb-Douglas aggregator). We now generalize this condition by assuming the demand condition takes the form:

$$p^\theta Y = E$$

In an aggregate context,  $\theta$  serves as a reduced-form parameter capturing the combined income and substitution effects in labor supply.<sup>20</sup> In an industry or small open economy context,  $\theta$  represents the demand elasticity.

Under this more general demand structure, the Value of Organizational Learning

<sup>20</sup> For example, if labor supply is  $N = p^{-\phi} Y^{-\rho}$ , where  $\phi$  is the Frisch elasticity and  $\rho$  the income elasticity, then, using the aggregate labor demand condition  $pY = \nu N$ , we obtain  $p^\theta Y = E$  with  $\theta \equiv \frac{1+\phi}{1+\rho}$ .

Technologies retains a sufficient statistic form:<sup>21</sup>

$$\text{VOLT}_\tau = (y_\tau \ell_\tau)^{\frac{\theta(1-\nu)}{\theta(1-\nu)+\nu}}$$

Thus, the exponent  $\frac{\theta(1-\nu)}{\theta(1-\nu)+\nu}$  captures how general equilibrium effects dampen or amplify the output gains from organizational learning. In our baseline case, price declines perfectly offset productivity gains to keep expenditure and employment constant, so the exponent simplifies to the decreasing returns to scale  $1 - \nu$  to organizational capital.

When demand is elastic ( $\theta > 1$ ), the price falls less than proportionately to output, causing total revenue and employment to rise. As demand becomes perfectly elastic, prices remain fixed, the exponent  $\frac{\theta(1-\nu)}{\theta(1-\nu)+\nu}$  approaches 1, and the aggregate output gain is linear in the efficiency gains  $y_\tau \ell_\tau$ .

When demand is inelastic ( $\theta < 1$ ), the increase in aggregate supply drives prices down so sharply that total revenue shrinks. Consequently, output goes up at the same time as employment goes down. As demand becomes perfectly inelastic,  $\frac{\theta(1-\nu)}{\theta(1-\nu)+\nu} \rightarrow 0$ , the price drop completely offsets the productivity gains, leaving aggregate output effectively unchanged.

This generalization highlights that our baseline estimates ( $\theta = 1$ ) are conservative if organizational learning technologies are applied to an aggregate economy with high labor supply elasticity, or industries and small open economies with high demand substitutability. For example, for an aggregate economy with labor supply  $N = p^{-\phi} Y^{-\rho}$ , suppose that the Frisch elasticity  $\phi$  is 0.5 (Chetty et al., 2011) and that there are no income effects (i.e.,  $\rho = 0$ ). This implies  $\theta = 1.5$  so  $\text{VOLT}_{46}$  goes from 2.08 to 2.66. At the industry level, reasonable estimates of  $\theta$  lie between 2 and 8 (Simonovska and Waugh, 2014; Boehm et al., 2023), suggesting that the median  $\text{VOLT}_{46}$  would be in the range of 2.79 and 6.47, rather than the median of 1.90 estimated in our baseline setting.

## 6 Discussion

This section places our findings in a broader perspective, connecting VOLT to debates on transformative AI, the nature of organizational learning, the evolution of business dynamism, and the productivity and organizational effects of AI adoption.

<sup>21</sup> To see this, we can start from the aggregate production function  $Y = (ZM)^{1-\nu} N^\nu$ . In equilibrium, employment is a fraction  $\nu$  of revenue,  $N = \nu p Y$ . Substituting this into the production function yields  $Y = ZM (\nu p)^{\frac{\nu}{1-\nu}}$ . Using the inverse demand function  $p = E^{\frac{1}{\theta}} Y^{-\frac{1}{\theta}}$  we get  $Y \propto (ZM)^{\frac{\theta(1-\nu)}{\theta(1-\nu)+\nu}}$ . Recalling that  $\text{VOLT} = Y_\tau^c / Y$  and noting that  $\frac{Z_\tau^c M_\tau^c}{ZM} = y_\tau \ell_\tau$ , we obtain the result.

## 6.1 Transformative AI

Current debates about AI’s transformative potential (e.g., [Agrawal et al., 2025](#)) have largely focused on two distinct channels. A growing body of work — including [Aghion et al. \(2017\)](#), [Aghion and Bunel \(2024\)](#), [Filippucci et al. \(2024\)](#) and [Acemoglu \(2025\)](#) — emphasizes automation of production tasks. Quantitatively, this literature finds modest effects on productivity growth, ranging from around 0.05–0.1 percentage points per year in conservative calibrations to about 0.25–0.6 percentage points per year in more optimistic scenarios. The central reason is that automation affects only a subset of tasks, adoption is gradual, and complementarities across tasks limit how much task-level cost savings propagate to aggregate output.

A parallel line of work asks whether AI’s potential to accelerate scientific discovery and innovation could instead generate persistent increases in growth rates (e.g., [Aghion et al., 2017](#); [Trammell and Korinek, 2025](#); [Jones, 2025](#); [Restrepo, 2025](#); [Davidson et al., 2026](#); [Jones and Tonetti, 2026](#)). This literature builds on task-based and semi-endogenous growth models in which AI enters the ideas production function.

[Jones \(2025\)](#) shows that the growth impact of AI-driven research acceleration depends critically on three forces: the share of research tasks AI can perform, AI’s relative productivity at those tasks, and the severity of bottlenecks in the remaining human tasks. Even very large gains in some research activities can yield only moderate aggregate effects if humans remain as bottlenecks.

[Restrepo \(2025\)](#) studies an extreme benchmark in which Artificial General Intelligence makes all economically relevant work reproducible using compute. In the long run, growth is then driven by the accumulation of computational resources. While this framework allows for sustained exponential growth, it does not generically imply explosive or singular dynamics, as growth rates are tied to the pace of compute accumulation.

[Jones and Tonetti \(2026\)](#) embed automation of both goods and idea production into a calibrated task-based growth model with strong complementarities. Although the model formally admits growth acceleration and even explosive paths, their benchmark calibration implies a slow transition: output is only about 4% higher by 2040 and roughly 19% higher by 2060 relative to a constant-growth counterfactual. The key mechanism is the presence of weak links — tasks that are hard to automate and constrain aggregate output even when most other tasks experience rapid productivity growth.<sup>22</sup>

Our work abstracts from both direct production automation and AI-driven accelera-

---

<sup>22</sup> Synthesizing this literature, [Jones \(2026\)](#) emphasizes that both “business-as-usual” and “accelerated growth” trajectories are plausible, but that historically grounded models tend to predict long and gradual transitions rather than abrupt macroeconomic breaks.

tion of frontier innovation and instead focuses on a complementary channel: AI as an organizational learning technology. We ask how technologies that accelerate organizational learning affect aggregate outcomes. Rather than altering the production of goods or ideas, these technologies compress the time it takes organizations to learn about their productivity, accumulate organizational capital, and sort through selection. In particular, they enable firms to learn the same information as before, but at a faster pace. We show that this channel can generate relatively large level effects in aggregate output.

## 6.2 Organizational Learning

Our framework conceptualizes organizational learning as a dual process. Firms accumulate organizational capital, which raises productivity with age, and simultaneously learn about their underlying fundamentals, which governs entry and exit. Much of the existing literature studies these channels in isolation, raising the natural question of whether their respective contributions to VOLT can be disentangled.

A tempting counterfactual would accelerate organizational capital accumulation without simultaneously accelerating learning about fundamentals. However, such a counterfactual is not empirically disciplined. While we observe the productivity trajectories of firms that survive to maturity, we do not observe the counterfactual paths of firms that exit early, nor how faster learning would have altered their exit decisions.

Despite this difficulty, we show that under a mild assumption on productivity dynamics, it is possible to isolate the value of the selection channel alone.

### 6.2.1 Isolating the Selection Channel

To isolate the role of selection, we construct a counterfactual in which learning about fundamentals is accelerated as in the baseline counterfactual, while the accumulation of organizational capital is left unchanged. In this environment, potential entrants perfectly anticipate which firms would exit before maturity in the factual economy. As a result, only survivor types enter, but they follow their factual productivity trajectories.

We define the value of faster selection as the ratio of stationary aggregate output in this selection-only economy to that of the factual economy. While we cannot pinpoint the exact magnitude of this value without additional assumptions about the unobserved productivity of exitors, we can derive a lower bound based on the sufficient statistics identified in Section 3. In this counterfactual economy, the mass of active firms still increases by the factor  $\ell_\tau$ . If we assume that survivors are, on average, at least as productive as exitors while young, the average productivity in this counterfactual economy is at least

as large as in the baseline counterfactual economy. Hence, the lifespan effect identified in Section 3 provides a lower bound on the value of faster selection.

This result implies that the large contribution of the lifespan statistic in our baseline results provides a lower bound on the economic value of screening viable business models. However, this does not imply that the organizational capital channel in isolation would be small, because the learning channels are not additive. In particular, accelerating organizational capital accumulation in isolation could endogenously reduce the exit rate and hence generate a lifespan effect as well.

### 6.2.2 Pre-entry vs post-entry learning about fundamentals

As detailed in Section 3, we formalize a technology that accelerates learning about fundamentals in two steps. First, the technology accelerates learning after entry, meaning a firm learns upon entry all the information which would have taken them  $\tau$  periods to learn in the factual economy. Second, the technology accelerates learning before entry, allowing potential entrants to foresee which firms exit before reaching maturity.

We cannot empirically implement a counterfactual that includes only the first step without the second. Evaluating such a scenario would require observing the exit decisions firms would have made under a different information set, which is not recoverable from life-cycle data. By contrast, our baseline counterfactual relies only on the factual behavior of survivors, allowing us to remain agnostic about these unobservable dynamics.

### 6.2.3 Payments to organizational capital

A related literature attempts to measure both the value of organizational capital to firm owners (e.g., [Eisfeldt and Papanikolaou, 2013](#)) and the payments to that capital ([Atkeson and Kehoe, 2005](#)). We can use our framework to ask how faster organizational learning would change the income share accruing to firm owners through payments to organizational capital.

In our model, *gross* payments associated with organizational capital are a constant share of aggregate income: variable profits sum to  $(1 - \nu) pY$  in both the factual and counterfactual economy. The object of interest, however, is how much of this gross flow is ultimately received by owners *net* of fixed operating costs ([Atkeson and Kehoe, 2005](#)).

Letting  $\alpha_o$  and  $\alpha_o^c$  denote the share of aggregate income paid to owners net organiza-

tional rents in the factual and counterfactual, respectively, we get:<sup>23</sup>

$$\alpha_o^c - \alpha_o = (1 - \nu - \alpha_o) \left(1 - \frac{1}{y_\tau}\right)$$

Two implications follow immediately. First, the increase in payments to organizational capital is higher the higher is mature productivity relative to average productivity ( $y_\tau$ ), and is bounded above by  $(1 - \nu) \left(1 - \frac{1}{y_\tau}\right)$ . Second — and importantly for interpretation — this increase does not depend on the relative mature lifespan  $\ell_\tau$ : faster learning changes the owners' rent share through the fixed-cost burden per unit of aggregate income, not through the stationary mass of firms.

In the context of manufacturing, [Atkeson and Kehoe \(2005\)](#) estimate that payments to owners of organization capital are about 3.3% of output in their benchmark specification. Taking  $\alpha_o = 0.033$ ,  $\nu = 0.75$ , and our baseline  $y$  (e.g.,  $y \approx 2.69$  at  $\tau = 46$ ), we obtain

$$\alpha_o^c - \alpha_o \approx (1 - 0.75 - 0.033) \left(1 - \frac{1}{2.69}\right) \approx 0.14,$$

so faster organizational learning can raise the share of aggregate income accruing to firm owners by roughly 14 percentage points under this benchmark.

### 6.3 Implications for Business Dynamism and Size Distribution

A substantial body of work identifies firm churn — the simultaneous entry and exit of businesses — as a fundamental driver of aggregate productivity and economic renewal. In theoretical models of endogenous growth, this process takes the form of creative destruction, whereby innovation by entrants renders incumbents obsolete and exit reallocates resources toward more productive uses ([Aghion and Howitt, 1990](#)). Empirically, this reallocation is a primary source of efficiency gains ([Foster et al., 2001](#)), and low or declining exit rates are often interpreted as evidence of economic sclerosis ([Decker et al., 2014](#)).

By abstracting from creative destruction, our framework offers a complementary perspective. Entrants draw productivity from a stationary distribution rather than displacing incumbents with superior vintages. In this setting, high exit rates — especially the disproportionately high failure of young firms — reflect poor ex ante project selection rather

---

<sup>23</sup> Since  $\int_0^M \pi_i di = (1 - \nu) pY$ , we can write  $\alpha_o^c = (1 - \nu) - \frac{f^c M^c}{p^c Y^c}$ . Hence,  $\alpha_o^c - \alpha_o = \frac{\int_0^M (\pi_i^c - f^c)}{p^c Y^c} - \frac{\int_0^M (\pi_i - f)}{pY} = -\frac{f^c M^c}{p^c Y^c} + \frac{fM}{pY}$ . The expression then follows from noting that  $\frac{f^c M^c}{p^c Y^c} = \frac{1}{y} \frac{fM}{pY}$ .

than efficient reallocation.

Through this lens, AI reduces churn by improving learning about fundamentals prior to entry, shifting selection to the pre-entry stage. As a result, aggregate output and productivity rise even as entry and exit rates fall. Improved organizational learning ensures that only firms with long-run viability enter the market, eliminating experimental churn. Although declining business dynamism is often interpreted as stagnation, here it represents an efficiency gain.

This increase in firm longevity also reduces average firm size. Because surviving firms live substantially longer than the typical entrant, the stationary equilibrium features a larger mass of active firms. With aggregate labor supply (or aggregate expenditure) fixed, these firms compete for the same pool of resources, generating a crowding effect that forces each firm to operate at a smaller scale. The firm size distribution therefore shifts from one characterized by many small, young firms and a few large incumbents to a more compressed distribution of smaller firms.

Finally, the technology compresses the stationary productivity distribution. In the baseline economy, dispersion reflects the coexistence of young firms climbing the learning curve and mature firms at their peak. Accelerated learning truncates the left tail by screening out low-quality potential entrants and immediately shifting the remaining mass to mature productivity levels. The result is a more homogeneous cross-section with substantially lower firm-level heterogeneity.

## 6.4 Productivity and Organizational Effects of AI Adoption

A growing empirical literature measures the productivity impact of generative AI in organizational settings, typically by randomizing access to an AI tool or exploiting quasi-experimental variation in rollout. These studies offer credible estimates of task-level and workflow-level effects, and they help clarify the organizational mechanisms through which AI may accelerate learning.

A recurring finding is that AI can increase measured productivity in knowledge intensive jobs, with gains that are often larger for less experienced workers. In a large field setting, [Brynjolfsson et al. \(2025\)](#) study the staggered deployment of a conversational assistant for customer-support agents and find an average increase of about 14–15 percent in issues resolved per hour, with substantially larger gains for novice agents and smaller gains for highly experienced agents. In professional writing tasks, [Noy and Zhang \(2023\)](#) find that access to ChatGPT reduces completion time by roughly 40 percent while increasing evaluated quality by about 18 percent, compressing the performance distribution.

In high-end knowledge work, Dell’Acqua et al. (2023) show that performance improvements can be large for tasks that lie within the model’s capability set, but that outcomes can deteriorate on tasks outside that set, highlighting the importance of verification costs, appropriate task selection, and organizational guardrails.

Evidence in software development points to potentially large effects but also substantial dependence on context and workflow integration. Laboratory-style tasks suggest sizable speed gains (e.g., Peng et al., 2023), and natural experiments around the availability of coding assistants across languages or repositories find increases in code output measures (e.g., Yeverehyahu et al., 2024). At the same time, recent evidence emphasizes that benefits are not guaranteed in realistic environments. Becker et al. (2025) study experienced open-source developers working on tasks in familiar repositories and find that allowing AI tools can slow completion time on average, alongside a large perception--reality gap.

A second line of micro evidence concerns collaboration and the internal organization of knowledge. Dell’Acqua et al. (2025) show that access to AI can substitute for some of the coordination benefits of teamwork in product-development tasks: individuals with AI assistance can approach the performance of human teams, and solutions become more balanced across functional dimensions. One interpretation is that AI reduces reliance on functional silos by allowing workers to bridge knowledge gaps and reuse dispersed organizational information more effectively. This is naturally connected to organizational learning: the technology can speed up the process by which workers and teams acquire, combine, and operationalize knowledge inside the firm.

Our paper complements this literature by asking how an improvement in the organizational learning technology translate into aggregate outcomes once filtered through firm life-cycle dynamics and general equilibrium. In particular, in our framework, faster organizational learning affects aggregate output not only through productivity but also by increasing average lifespans and hence the stationary mass of active firms.

Micro estimates are naturally interpreted as fixed-price, partial-equilibrium effects; once adoption is broad, prices and quantities adjust, and the mapping from local efficiency gains to aggregate output becomes concave. To illustrate the magnitude of this equilibrium compression, consider our baseline estimates at  $\tau = 46$ , for which  $y \approx 2.69$  and  $\ell \approx 7$ , implying  $y\ell \approx 18.8$ . Under a fixed-price (partial-equilibrium) experiment, aggregate output scales linearly with this efficiency term, so the implied sector-wide effect is proportional to  $y\ell$ . By contrast, in general equilibrium with span-of-control  $\nu = 0.75$ , aggregate output responds according to  $\text{VOLT} = (y\ell)^{1-\nu} = (18.8)^{0.25} \approx 2.08$ . Thus, equilibrium price adjustment compresses the fixed-price effect by an order of magnitude. This

gap between fixed-price and equilibrium effects cautions against mechanically extrapolating micro productivity gains to economy-wide implications of AI adoption.

## 7 Conclusion

This paper takes a first step toward measuring the value of AI as an organizational learning technology. We show that in a broad class of firm dynamics models, the value of organizational learning technologies (VOLT) is determined by two simple statistics: the relative size and lifespan of mature firms. This sufficient statistic result provides a transparent way to assess the aggregate output gains from faster organizational learning without relying on specific assumptions about firm learning.

Using U.S. data, we find that VOLT is quantitatively large — on the order of one GDP — implying that accelerating organizational learning could double aggregate output. Across industries, VOLT varies widely and is largely orthogonal to existing measures of AI exposure based on adoption feasibility, indicating that it captures a distinct dimension of AI’s economic impact.

A central lesson of this paper is that the importance of the lifespan channel — and, more generally, how faster learning maps into aggregate outcomes — depends on the market structure governing entry and the evolution of the firm population. In our benchmark class of models, the flow of entrants is fixed, so faster learning raises output partly by extending firm lifespans and increasing the mass of active firms. At the same time, faster learning increases the expected value of entry by reducing fixed-cost burdens and front-loading profits. If entry were elastic rather than fixed, this higher entry value would translate into a higher entry flow and an even larger mass of active firms, implying that our fixed-entry measure of VOLT is conservative.

Conversely, in an environment where the mass of firms is fixed by scarce managerial capacity, the lifespan effect is neutralized and the gains from faster learning operate exclusively through productivity and selection. Pinning down which entry margin is empirically relevant is therefore central not only for quantifying the gains from organizational learning acceleration, but also for evaluating the broader class of shocks and policies that affect firm productivity and longevity. More broadly, this perspective shifts attention from within-firm productivity effects toward life-cycle selection: AI may improve screening and selection by shifting learning to the pre-entry stage, reducing premature exit among young firms, and changing survival patterns across sectors. Measuring these life-cycle margins is therefore an important direction for future work.

More generally, our results suggest that assessments of AI’s impact should extend

beyond the within-organization productivity effects emphasized in much of the micro evidence. While task-level treatment effects are informative about mechanisms, our estimates indicate that their aggregate implications can be an order of magnitude smaller once general equilibrium adjustments are taken into account.

# Appendix

## A The Organizational Learning Technology Increases the Net Present Value of Entry

In this section, we assume that average productivity among young firms is weakly increasing with age, and that the expected average productivity of a firm after it matures is greater than its average productivity while young.

$$\bar{z}_k \leq \bar{z}_{k'} \text{ for all } k \leq k' \leq \tau, \text{ and } \zeta_\tau(S) \geq \frac{1}{n} \sum_{k=1}^{\tau} z_k^{\frac{1}{1-\nu}}(S) \quad (\text{A.1})$$

Let  $A_k$  be 1 if the entrant is alive at age  $k$  in the factual and 0 otherwise. Let  $a_k := \mathbb{E}_\mu [A_k \pi_k]$  be the pre-entry expected age- $k$  contribution to the undiscounted lifetime value  $W := \sum_{k \geq 0} a_k$ , so the discounted net present value is  $V(\beta) := \sum_{k \geq 0} \beta^k a_k$ . For the accelerated learning counterfactual, define  $A_k^c, a_k^c, W^c$  and  $V^c$  analogously, and define:

$$\Delta_k \equiv a_k^c - a_k, \Delta W \equiv W^c - W = \sum_{k \geq 0} \Delta_k \text{ and } \Delta V(\beta) \equiv V^c(\beta) - V(\beta) = \sum_{k \geq 0} \beta^k \Delta_k$$

We start by showing that  $\Delta_k \geq 0$  for all  $k \leq \tau$ . By construction of the learning counterfactual, entrants have weakly higher survival probability to each young age  $k \leq \tau$  in the counterfactual than in the factual. Also, (A.1) implies that entrants have weakly higher expected productivity at each young age in the counterfactual than in the factual. Therefore,  $a_k^c \geq a_k$  for every  $k \leq \tau$ .

Next, we show that there exists a constant  $\kappa > 0$  such that  $a_k^c = \kappa a_k$  for all  $k > \tau$ . Consequently, either  $\Delta_k \geq 0$  for all  $k > \tau$  (if  $\kappa \geq 1$ ) or  $\Delta_k \leq 0$  for all  $k > \tau$  (if  $\kappa \leq 1$ ). For  $k > \tau$ , the counterfactual does not change firms' productivity. Moreover, as we argued in the proof of Proposition 1 in Section 3.2, the equilibrium price change scales per-period profits proportionally and leaves exit decisions unchanged. Combining these two facts implies that the entire sequence of expected old-age contributions  $\{a_k^c\}_{k > \tau}$  differs from  $\{a_k\}_{k > \tau}$  only by a common multiplicative factor  $\kappa > 0$ , i.e.  $a_k^c = \kappa a_k$  for all  $k > \tau$ .

**Lemma 1.** *If  $\Delta W \geq 0$ , then  $\Delta V(\beta) \geq 0$  for every  $\beta \in (0, 1)$ .*

*Proof.* If  $\kappa \geq 1$ , then  $\Delta_k \geq 0$  for all  $k$ , so  $\Delta V(\beta) = \sum_{k \geq 0} \beta^k \Delta_k \geq 0$  for all  $\beta \in (0, 1)$ . Suppose instead  $\kappa < 1$ . In this case, we have  $\Delta_k \geq 0$  for  $k \leq \tau$  but  $\Delta_k \leq 0$  for  $k > \tau$ . Let  $C_k \equiv \sum_{j=0}^k \Delta_j$ . The sequence  $\{C_k\}$  is weakly increasing on  $\{0, \dots, \tau\}$  and weakly

decreasing for  $k \geq \tau$ . Moreover,

$$\lim_{k \rightarrow \infty} C_k = \sum_{j \geq 0} \Delta_j = \Delta W \geq 0.$$

It follows that  $C_k \geq 0$  for all  $k \geq 0$ . Finally, for any  $\beta \in (0, 1)$ , we have:

$$\Delta V(\beta) = \sum_{k \geq 0} \beta^k \Delta_k = \sum_{k \geq 0} \beta^k (C_k - C_{k-1}) = (1 - \beta) \sum_{k \geq 0} \beta^k C_k$$

Since  $1 - \beta > 0$ ,  $\beta^k > 0$ , and  $C_k \geq 0$  for all  $k$ , the right-hand side is nonnegative. Hence  $\Delta V(\beta) \geq 0$  for every  $\beta \in (0, 1)$ , completing the proof.  $\square$

## B Aggregate Sufficient Statistics and VOLT

**Table 1:** Aggregate Results by Maturity Threshold; Baseline  $\nu = 0.75$ .

Threshold ( $\tau$ )	VOLT	$\ell_\tau$	$y_\tau$	$\bar{\ell}_{>\tau}$	$\bar{n}_{>\tau}$	$\frac{\ell_\tau^{1-\nu}-1}{VOLT-1}$	$\frac{y_\tau^{1-\nu}-1}{VOLT-1}$
0	1.01	1.00	1.05	9.8	19.3	0.00	1.00
1	1.09	1.27	1.10	12.4	20.1	0.71	0.28
2	1.13	1.46	1.13	14.3	20.8	0.74	0.24
3	1.17	1.62	1.16	15.8	21.4	0.75	0.23
4	1.21	1.77	1.20	17.3	21.9	0.74	0.22
5	1.24	1.92	1.22	18.7	22.5	0.74	0.22
10	1.37	2.60	1.35	25.4	24.8	0.73	0.21
15	1.47	3.22	1.46	31.5	26.8	0.72	0.21
20	1.57	3.82	1.59	37.3	29.2	0.70	0.22
25	1.65	4.39	1.70	42.9	31.2	0.68	0.22
46	2.08	7.00	2.69	68.4	49.4	0.58	0.26

**Table 2:** Aggregate Results by Maturity Threshold;  $\nu = .7$ 

Threshold ( $\tau$ )	VOLT	$\ell_\tau$	$y_\tau$	$\frac{\ell_\tau^{1-\nu}-1}{\text{VOLT}-1}$	$\frac{y_\tau^{1-\nu}-1}{\text{VOLT}-1}$
0	1.02	1.00	1.05	0.00	1.00
1	1.10	1.27	1.10	0.71	0.27
2	1.16	1.46	1.13	0.74	0.24
3	1.21	1.62	1.16	0.74	0.22
4	1.25	1.77	1.20	0.74	0.22
5	1.29	1.92	1.22	0.74	0.21
10	1.46	2.60	1.35	0.72	0.21
15	1.59	3.22	1.46	0.71	0.20
20	1.72	3.82	1.59	0.69	0.21
25	1.83	4.39	1.70	0.67	0.21
46	2.41	7.00	2.69	0.56	0.24

**Table 3:** Aggregate Results by Maturity Threshold;  $\nu = .8$ 

Threshold ( $\tau$ )	VOLT	$\ell_\tau$	$y_\tau$	$\frac{\ell_\tau^{1-\nu}-1}{\text{VOLT}-1}$	$\frac{y_\tau^{1-\nu}-1}{\text{VOLT}-1}$
0	1.01	1.00	1.05	0.00	1.00
1	1.07	1.27	1.10	0.71	0.28
2	1.11	1.46	1.13	0.74	0.24
3	1.14	1.62	1.16	0.75	0.23
4	1.16	1.77	1.20	0.75	0.23
5	1.19	1.92	1.22	0.75	0.22
10	1.29	2.60	1.35	0.74	0.22
15	1.36	3.22	1.46	0.73	0.22
20	1.43	3.82	1.59	0.71	0.22
25	1.50	4.39	1.70	0.69	0.23
46	1.80	7.00	2.69	0.60	0.27

## C Industry Level VOLT and LLM Exposure

**Table 4:** Industry-level estimates of VOLT, its components, and LLM exposure.

Code	Industry	VOLT <sub>46</sub>	$\ell_{46}$	$y_{46}$	$\ell^{1-\nu}$	$y^{1-\nu}$	LLM Exposure
113	Forestry & Logging	1.93	8.07	1.73	1.69	1.15	0.23
115	Ag Support	2.07	8.10	2.28	1.69	1.23	0.18
211	Oil & Gas Extraction	1.81	6.53	1.64	1.60	1.13	0.33
212	Mining (excl Oil/Gas)	1.62	4.53	1.50	1.46	1.11	0.20
213	Mining Support	2.15	5.13	4.13	1.50	1.43	0.23
221	Utilities	1.74	5.40	1.68	1.52	1.14	0.35
236	Construction of Bldgs	2.49	9.49	4.07	1.76	1.42	0.31
237	Heavy/Civil Eng Const	1.94	6.62	2.13	1.60	1.21	0.22
238	Specialty Trade Const	2.15	7.36	2.92	1.65	1.31	0.22
311	Food Mfg	2.06	6.42	2.82	1.59	1.30	0.19
312	Beverage & Tobacco	2.45	4.90	7.35	1.49	1.65	0.25
313	Textile Mills	1.94	7.72	1.82	1.67	1.16	0.23
314	Textile Prod Mills	2.00	7.25	2.19	1.64	1.22	0.23
315	Apparel Mfg	2.88	21.81	3.17	2.16	1.33	0.26
316	Leather & Allied	2.06	6.98	2.61	1.63	1.27	0.18
321	Wood Prod Mfg	1.83	5.39	2.10	1.52	1.20	0.19
322	Paper Mfg	1.55	3.92	1.49	1.41	1.10	0.20
323	Printing	1.79	5.38	1.92	1.52	1.18	0.35
324	Petroleum & Coal	1.79	3.31	3.07	1.35	1.32	0.30
325	Chemical Mfg	2.02	8.20	2.05	1.69	1.20	0.32
326	Plastics & Rubber	1.61	4.05	1.67	1.42	1.14	0.23
327	Nonmetallic Mineral	1.63	3.49	2.01	1.37	1.19	0.25
331	Primary Metal	1.51	3.51	1.50	1.37	1.11	0.23
332	Fabricated Metal	1.72	4.91	1.77	1.49	1.15	0.27
333	Machinery Mfg	1.78	5.25	1.89	1.51	1.17	0.31
334	Computer/Electronic	1.92	6.49	2.11	1.60	1.21	0.45
335	Elec Equip & Appl	1.93	5.46	2.52	1.53	1.26	0.29
336	Transportation Equip	1.96	6.00	2.45	1.57	1.25	0.28
337	Furniture Mfg	2.10	6.71	2.90	1.61	1.30	0.24
339	Misc Mfg	1.98	6.23	2.45	1.58	1.25	0.31
423	Wholesale: Durable	1.85	6.93	1.67	1.62	1.14	0.41
424	Wholesale: Nondurable	1.90	6.15	2.10	1.57	1.20	0.38
425	Wholesale: Electronic	1.98	7.85	1.97	1.67	1.19	0.49
441	Motor Vehicle Dealers	1.66	4.39	1.74	1.45	1.15	0.30
442	Electronics Stores	1.62	6.83	1.01	1.62	1.00	0.39
442	Furniture/Home Stores	1.48	4.43	1.07	1.45	1.02	0.39

*Continued on next page*

Code	Industry	VOLT <sub>46</sub>	$\ell_{46}$	$y_{46}$	$\ell^{1-\nu}$	$y^{1-\nu}$	LLM Exposure
444	Bldg Material Dealers	1.27	3.87	0.67	1.40	0.90	0.32
445	Food & Bev Stores	1.78	5.72	1.75	1.55	1.15	0.25
446	Health/Personal Care	1.43	3.83	1.10	1.40	1.03	0.38
447	Gasoline Stations	1.45	4.12	1.07	1.42	1.02	0.26
448	Clothing Stores	1.53	4.21	1.30	1.43	1.07	0.36
451	Misc Store Retailers	1.60	5.74	1.15	1.55	1.04	0.35
451	Nonstore Retailers	1.90	11.05	1.19	1.82	1.04	0.35
451	Sport/Hobby/Book	1.57	5.10	1.18	1.50	1.04	0.35
452	General Merchandise	1.36	1.87	1.85	1.17	1.17	0.23
481	Air Transportation	2.43	6.74	5.20	1.61	1.51	0.29
483	Water Transportation	2.24	6.45	3.91	1.59	1.41	0.30
484	Truck Transportation	2.63	13.56	3.51	1.92	1.37	0.31
485	Transit & Ground	2.27	9.75	2.70	1.77	1.28	0.31
487	Scenic/Sightseeing	2.17	8.89	2.50	1.73	1.26	0.29
488	Support for Trans	1.90	7.11	1.84	1.63	1.16	0.35
492	Couriers & Msgrs	3.41	13.97	9.63	1.93	1.76	0.25
493	Warehousing	1.43	4.21	0.98	1.43	1.00	0.20
511	Publishing	1.79	7.86	1.30	1.67	1.07	0.63
512	Motion Picture/Sound	2.44	11.29	3.13	1.83	1.33	0.39
515	Broadcasting	2.13	8.76	2.33	1.72	1.24	0.58
517	Telecommunications	1.79	9.70	1.06	1.76	1.01	0.47
518	Data Processing	2.18	8.57	2.62	1.71	1.27	0.63
519	Other Info Svcs	1.94	18.13	0.79	2.06	0.94	0.63
522	Credit Intermediation	1.74	7.93	1.16	1.68	1.04	0.57
523	Securities/Finc	2.36	6.05	5.13	1.57	1.51	0.56
524	Insurance Carriers	2.13	6.93	2.99	1.62	1.32	0.58
531	Real Estate	1.92	7.83	1.72	1.67	1.15	0.39
532	Rental & Leasing	1.74	4.53	2.05	1.46	1.20	0.35
533	Lessors Nonfin Assets	2.61	8.93	5.19	1.73	1.51	0.51
541	Prof/Sci/Tech Svcs	2.15	7.46	2.85	1.65	1.30	0.55
551	Mgmt of Companies	2.07	8.38	2.17	1.70	1.21	0.53
561	Admin & Support	1.84	8.31	1.38	1.70	1.08	0.28
562	Waste Mgmt	1.86	5.70	2.11	1.55	1.21	0.27
611	Educational Svcs	2.65	8.53	5.78	1.71	1.55	0.38
621	Ambulatory Health	1.72	4.59	1.89	1.46	1.17	0.37
622	Hospitals	1.51	3.57	1.46	1.37	1.10	0.34
623	Nursing/Resid Care	2.11	7.13	2.80	1.63	1.29	0.22
624	Social Assistance	1.88	6.45	1.95	1.59	1.18	0.24
711	Performing Arts	2.87	11.02	6.15	1.82	1.58	0.35
712	Museums/Historical	2.02	6.29	2.66	1.58	1.28	0.37
713	Amusement/Gambling	1.95	7.35	1.97	1.65	1.18	0.24

*Continued on next page*

Code	Industry	VOLT <sub>46</sub>	$\ell_{46}$	$y_{46}$	$\ell^{1-\nu}$	$y^{1-\nu}$	LLM Exposure
721	Accommodation	1.74	5.95	1.55	1.56	1.12	0.23
722	Food Svcs & Drinking	1.65	5.85	1.25	1.55	1.06	0.17
811	Repair & Maint	1.67	6.24	1.25	1.58	1.06	0.24
812	Personal/Laundry	1.87	6.78	1.79	1.61	1.16	0.23
813	Religious/Civic Org	1.75	5.72	1.63	1.55	1.13	0.42

## References

- Daron Acemoglu. The simple macroeconomics of AI. *Economic Policy*, 40(121):13–58, 2025.
- Philippe Aghion and Simon Bunel. AI and growth: where do we stand? *Policy Note*, 2024.
- Philippe Aghion and Peter Howitt. A model of growth through creative destruction, 1990.
- Philippe Aghion, Benjamin F Jones, and Charles I Jones. Artificial intelligence and economic growth. Technical report, National Bureau of Economic Research, 2017.
- Ajay Agrawal, Joshua Gans, and Avi Goldfarb. *Prediction machines, updated and expanded: The simple economics of artificial intelligence*. Harvard Business Press, 2022.
- Ajay Agrawal, Anton Korinek, and Erik Brynjolfsson, editors. *The Economics of Transformative AI*. University of Chicago Press, Chicago, 2025.
- Ufuk Akcigit, Salomé Baslandze, and Francesca Lotti. Connecting to power: political connections, innovation, and firm dynamics. *Econometrica*, 91(2):529–564, 2023.
- Dario Amodei. Machines of loving grace: How AI could transform the world for the better, 2024. October, 2024. Available at <https://darioamodei.com/machines-of-loving-grace> (accessed October 31, 2025).
- Ruth Appel, Peter McCrory, Alex Tamkin, Miles McCain, Tyler Neylon, and Michael Stern. Anthropic economic index: AI’s impact on the workforce. White paper, Anthropic, 2025.
- Linda Argote. *Organizational learning: Creating, retaining and transferring knowledge*. Springer Science & Business Media, 2012.
- Linda Argote, Sara L Beckman, and Dennis Epple. The persistence and transfer of learning in industrial settings. *Management Science*, 36(2):140–154, 1990.

- Costas Arkolakis, Theodore Papageorgiou, and Olga A Timoshenko. Firm learning and growth. *Review of Economic Dynamics*, 27:146–168, 2018.
- Kenneth J Arrow. The economic implications of learning by doing. *The Review of Economic Studies*, 29(3):155–173, 1962.
- Andrew Atkeson and Patrick J. Kehoe. Modeling and measuring organization capital. *Journal of Political Economy*, 113(5):1026–1053, 2005.
- Esther Bailey, Daniel Fehder, Eric J. Floyd, Yael V. Hochberg, and Daniel J. Lee. Learning to quit? a multi-year, multi-site field experiment with innovation-driven entrepreneurs. NBER Working Paper 34755, National Bureau of Economic Research, January 2026.
- Joel Becker, Nate Rush, Beth Barnes, and David Rein. Measuring the impact of early-2025 AI on experienced open-source developer productivity, 2025.
- Martin Beraja and Nathan Zorzi. Inefficient automation. *The Review of Economic Studies*, 92(1):69–96, 2025.
- Lauren Falcao Bergquist, Danial Lashkari, and Eric Verhoogen. Wedges: A microeconomic perspective on misallocation. Technical report, ROCKWOOL Foundation Berlin (RFBerlin), 2026.
- Nicholas Bloom, Raffaella Sadun, and John Van Reenen. *Management as a Technology?*, volume 22327. National Bureau of Economic Research Cambridge, MA, 2016.
- Christoph E Boehm, Andrei A Levchenko, and Nitya Pandalai-Nayar. The long and short (run) of trade elasticities. *American Economic Review*, 113(4):861–905, 2023.
- Erik Brynjolfsson, Tom Mitchell, and Daniel Rock. What can machines learn and what does it mean for occupations and the economy? In *AEA papers and proceedings*, volume 108, pages 43–47. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203, 2018.
- Erik Brynjolfsson, Danielle Li, and Lindsey Raymond. Generative AI at work. *The Quarterly Journal of Economics*, 140(2):889–942, 2025.
- Francisco J Buera and Yongseok Shin. Financial frictions and the persistence of history: A quantitative exploration. *Journal of Political Economy*, 121(2):221–272, 2013.
- Francisco J Buera, Joseph P Kaboski, and Yongseok Shin. Finance and development: A tale of two sectors. *The American Economic Review*, 101(5):1964–2002, 2011.

- Aaron Chatterji, Thomas Cunningham, David J Deming, Zoe Hitzig, Christopher Ong, Carl Yan Shan, and Kevin Wadman. How people use chatgpt. Technical report, National Bureau of Economic Research, 2025.
- Raj Chetty, Adam Guren, Day Manoli, and Andrea Weber. Are micro and macro labor supply elasticities consistent? a review of evidence on the intensive and extensive margins. *American Economic Review*, 101(3):471–475, 2011.
- Gian Luca Clementi and Hugo A Hopenhayn. A theory of financing constraints and firm dynamics. *The Quarterly Journal of Economics*, 121(1):229–265, 2006.
- Susan Cohen, Daniel C. Fehder, Yael V. Hochberg, and Fiona Murray. The design of startup accelerators. *Research Policy*, 48(7):1781–1797, 2019.
- Tom Davidson, Basil Halperin, Thomas Houlden, and Anton Korinek. When does automating AI research produce explosive growth? 2026.
- Jan De Loecker, Jan Eeckhout, and Gabriel Unger. The rise of market power and the macroeconomic implications. *The Quarterly Journal of Economics*, 135(2):561–644, 2020.
- Ryan Decker, John Haltiwanger, Ron Jarmin, and Javier Miranda. The role of entrepreneurship in us job creation and economic dynamism. *Journal of Economic Perspectives*, 28(3):3–24, 2014.
- Fabrizio Dell’Acqua, Edward McFowland, Ethan Mollick, Hila Lifshitz-Assaf, Katherine C. Kellogg, Saran Rajendran, Lisa Krayer, François Candelon, and Karim R. Lakhani. Navigating the jagged technological frontier: Field experimental evidence of the effects of AI on knowledge worker productivity and quality. *Harvard Business School Working Paper*, (24-013), September 2023. Available at SSRN: <https://ssrn.com/abstract=4573321>.
- Fabrizio Dell’Acqua, Charles Ayoubi, Hila Lifshitz, Raffaella Sadun, Ethan Mollick, Lilach Mollick, Yi Han, Jeff Goldman, Hari Nair, Stewart Taub, et al. The cybernetic teammate: A field experiment on generative AI reshaping teamwork and expertise. Technical report, National Bureau of Economic Research, 2025.
- Avinash K Dixit and Joseph E Stiglitz. Monopolistic competition and optimum product diversity. *The American Economic Review*, 67(3):297–308, 1977.
- Andrea L Eisfeldt and Dimitris Papanikolaou. Organization capital and the cross-section of expected returns. *The Journal of Finance*, 68(4):1365–1406, 2013.

- Tyna Eloundou, Sam Manning, Pamela Mishkin, and Daniel Rock. Gpts are gpts: Labor market impact potential of LLMs. *Science*, 384(6702):1306–1308, 2024.
- Francesco Filippucci, Peter Gal, and Matthias Schief. Miracle or myth? assessing the macroeconomic productivity gains from artificial intelligence. Technical report, OECD Publishing, 2024.
- Lucia Foster, John C Haltiwanger, and Cornell John Krizan. Aggregate productivity growth: Lessons from microeconomic evidence. In *New developments in productivity analysis*, pages 303–372. University of Chicago Press, 2001.
- Luis Garicano. Hierarchies and the organization of knowledge in production. *Journal of Political Economy*, 108(5):874–904, 2000.
- Robert M Grant. Toward a knowledge-based theory of the firm. *Strategic Management Journal*, 17(S2):109–122, 1996.
- Nezih Guner, Gustavo Ventura, and Yi Xu. Macroeconomic implications of size-dependent policies. *Review of Economic Dynamics*, 11(4):721–744, 2008.
- Hugo Hopenhayn and Richard Rogerson. Job turnover and policy evaluation: A general equilibrium analysis. *Journal of Political Economy*, 101(5):915–938, 1993.
- Hugo A Hopenhayn. Entry, exit, and firm dynamics in long run equilibrium. *Econometrica*, 60(5):1127–1150, 1992.
- Hugo A Hopenhayn. Firms, misallocation, and aggregate productivity: A review. *Annual Review of Economics*, 6(1):735–770, 2014.
- Chang-Tai Hsieh and Peter J Klenow. The life cycle of plants in India and Mexico. *The Quarterly Journal of Economics*, 129(3):1035–1084, 2014.
- Enrique Ide. Automation, AI, and the intergenerational transmission of knowledge. *arXiv preprint arXiv:2507.16078*, 2025.
- Enrique Ide and Eduard Talamàs. Artificial intelligence in the knowledge economy. *Journal of Political Economy*, 133(12), December 2025.
- Benjamin F Jones. Artificial intelligence in research and development. Technical report, National Bureau of Economic Research, 2025.

- Charles I. Jones. A.i. and our economic future. Manuscript, Stanford Graduate School of Business and NBER. In preparation for the *Journal of Economic Perspectives*, January 2026.
- Charles I. Jones and Christopher Tonetti. Past automation and future a.i.: How weak links tame the growth explosion. Very preliminary manuscript, Stanford Graduate School of Business and NBER, January 2026.
- Boyan Jovanovic. Selection and the evolution of industry. *Econometrica*, 50(3):649–670, 1982. doi: 10.2307/1912606.
- Benjamin Labaschin, Tyna Eloundou, Sam Manning, Pamela Mishkin, and Daniel Rock. Extending “gpts are gpts” to firms. In *AEA Papers and Proceedings*, volume 115, pages 51–55. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203, 2025.
- Santiago Levy. *Good intentions, bad outcomes: Social policy, informality, and economic growth in Mexico*. Rowman & Littlefield, 2010.
- Robert E. Lucas, Jr. On the size distribution of business firms. *The Bell Journal of Economics*, 9(2):508–523, 1978.
- Erzo GJ Luttmer. Selection, growth, and the size distribution of firms. *The Quarterly Journal of Economics*, 122(3):1103–1144, 2007.
- James G March. Exploration and exploitation in organizational learning. *Organization science*, 2(1):71–87, 1991.
- Marc J. Melitz. The impact of trade on intra-industry reallocations and aggregate industry productivity. *Econometrica*, 71(6):1695–1725, 2003.
- Virgiliu Midrigan and Daniel Yi Xu. Finance and misallocation: Evidence from plant-level data. *American Economic Review*, 104(2):422–458, 2014.
- Sendhil Mullainathan and Ashesh Rambachan. Science in the age of algorithms. *NBER Chapters*, 2025.
- Richard R Nelson and Sidney G Winter. *An evolutionary theory of economic change*. harvard university press, 1985.
- Ikujiro Nonaka. A dynamic theory of organizational knowledge creation. *Organization science*, 5(1):14–37, 1994.

- Shakked Noy and Whitney Zhang. Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654):187–192, 2023.
- Ariel Pakes and Richard Ericson. Empirical implications of alternative models of firm dynamics. *Journal of Economic Theory*, 79(1):1–45, 1998.
- Sida Peng, Eirini Kalliamvakou, Peter Cihon, and Mert Demirer. The impact of AI on developer productivity: Evidence from github copilot, 2023.
- Jesse Perla and Christopher Tonetti. Equilibrium imitation and growth. *Journal of Political Economy*, 122(1):52–76, 2014.
- Michael Polanyi. *The Tacit Dimension*. Doubleday & Co, 1966.
- Edward C Prescott and Michael Visscher. Organization capital. *Journal of political Economy*, 88(3):446–461, 1980.
- Pascual Restrepo. We won’t be missed: Work and growth in the era of AGI. Technical report, National Bureau of Economic Research, 2025.
- Ina Simonovska and Michael E Waugh. The elasticity of trade: Estimates and evidence. *Journal of International Economics*, 92(1):34–50, 2014.
- Kiran Tomlinson, Sonia Jaffe, Will Wang, Scott Counts, and Siddharth Suri. Working with AI: Measuring the occupational implications of generative AI. *arXiv preprint arXiv:2507.07935*, 2025.
- Philip Trammell and Anton Korinek. Economic growth under transformative AI. Working paper, National Bureau of Economic Research, 2025. Invited for the Annual Review of Economics, September 2025.
- Silvio Vismara, Gresa Latifi, Leonard Meinzinger, and Alexander Pass. Generative AI-powered venture screening: Can large language models help venture capitalists? *International Review of Financial Analysis*, page 104748, 2025.
- Doron Yeverechyahu, Raveesh Mayya, and Gal Oestreicher-Singer. The impact of large language models on open-source innovation: Evidence from github copilot, 2024.