

The Value of Organizational Learning Technologies*

Martin Beraja

UC Berkeley

Eduard Talamàs

IESE Barcelona

March 23rd, 2026

Abstract

Organizations learn over time. They build organizational capital and form beliefs about their fundamentals. Motivated by recent advances in AI, we study organizational learning technologies that accelerate both processes. We show that, in a large class of models of firm dynamics, 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 organizational learning technologies like AI have the potential to double aggregate output. Much of VOLT reflects increases in average firm lifespans rather than productivity. Across industries, VOLT varies widely and is orthogonal to existing AI exposure measures. Overall, our results point to faster organizational learning as a meaningful and distinct channel of AI’s transformative potential, beyond production automation and scientific discovery.

1 Introduction

Organizations learn over time. They build organizational capital, becoming more productive with age (Atkeson and Kehoe, 2005; Eisefeldt and Papanikolaou, 2013; Crouzet and Eberly, 2021), and form beliefs about their fundamentals, shaping their entry and exit decisions (Jovanovic, 1982). Through its effects on firm productivity and selection, slow organizational learning depresses aggregate productivity (Hsieh and Klenow, 2014)

*Martin Beraja: maberaja@berkeley.edu. Eduard Talamàs: etalamas@iese.edu. We are grateful to Chad Jones for his insightful discussion at the 2026 San Francisco Fed Growth Conference, and to participants across the various seminars and conferences where this paper has been presented for their valuable feedback. Eduard acknowledges the financial support of IESE through the High Impact Initiative-course 2024/2026. Karan Makkar provided excellent research assistance. All errors are our own.

and hinders the reallocation of resources across firms (Buera and Shin, 2013; Midrigan and Xu, 2014).

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., 2025a), and they help codify and diffuse tacit knowledge in organizations (Brynjolfsson et al., 2025). Indeed, AI-driven organizational memory systems — such as *Glean* and *o9's Digital Brain* — are already being adopted, helping firms build organizational capital faster. In addition, AI improves prediction (Agrawal et al., 2022), enabling faster learning about fundamentals. New AI prediction systems — such as *Aaru*, *Expected Parrot* and *Shopify's SimGym* — produce increasingly realistic simulations that firms are using to evaluate new ideas, and similar tools already outperform human analysts in predicting early-stage startup survival (Vismara et al., 2025).

In this paper, we conceptualize organizational learning technologies and study their aggregate implications. Specifically, in a broad class of models of firm dynamics à la Hopenhayn (1992), we consider technologies that accelerate organizational learning: they allow firms to (i) accumulate organizational capital faster, in the sense that young firms reach mature levels of productivity earlier, and (ii) learn about their fundamentals earlier, in the sense that potential entrants can predict which projects are good enough to survive to maturity.

We define the Value of Organizational Learning Technologies (VOLT) as the ratio of aggregate 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 ℓ . Formally,

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

where ν is the span of control parameter (Lucas, 1978) that governs decreasing returns to scale.¹ The size effect $y^{1-\nu}$ corresponds to the increase in Total Factor Productivity (TFP) and it is larger when mature firms are relatively more productive, either because organizational capital accumulates slowly or because survivors are positively selected on productivity. The lifespan effect $\ell^{1-\nu}$ reflects the increase in the number of firms and it is larger when young firms exit at higher rates, for example because entrants struggle to identify projects that would survive to maturity. Both effects are amplified as decreasing returns $1 - \nu$ become stronger.

¹ Alternatively, under monopolistic competition and CES preferences (Dixit and Stiglitz, 1977), the relevant exponent is $1/(\sigma - 1)$ instead of $1 - \nu$, where $\sigma > 1$ denotes the elasticity of substitution across varieties.

Using 2023 U.S. establishment-level data from the Business Dynamics Statistics (BDS), we estimate 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 establishment 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 output, while using a maturity threshold of 5 years implies a 22% increase. These results point to faster organizational learning as a meaningful and distinct channel of AI’s transformative potential, alongside production automation and scientific discovery which have received most attention to date.

Approximately one quarter of the increase in output is driven by higher TFP, while the remaining reflects a larger mass of firms. This finding indicates that greater firm longevity can be a powerful mechanism through which organizational learning technologies affect aggregate output, beyond increases in TFP that have been the focus of much of the literature.

Disaggregating VOLT across industries reveals substantial heterogeneity. In particular, we find that VOLTs are largely orthogonal to standard measures of LLM exposure (Eloundou et al., 2024) based on adoption feasibility. A rich off-diagonal pattern emerges: many industries where LLMs appear easy to adopt yield relatively modest gains, while many low-exposure industries harbor large gains from faster learning. These results suggest that what is easiest to automate with AI is not necessarily what is most valuable, and vice versa.

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

Firm Dynamics and Organizational Learning. A foundational literature views the process of organizational learning as an important driver of firm dynamics and aggregate outcomes (e.g., Jovanovic, 1982; Pakes and Ericson, 1998; Atkeson and Kehoe, 2005; Hsieh

and Klenow, 2014; Perla and Tonetti, 2014), with recent work studying the role of data (e.g., Baley and Veldkamp, 2025). We take these properties as given and quantify the macroeconomic value of the technologies — such as AI — that use data to accelerate organizational learning. 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 find that these payments are substantial: roughly 4 percent of manufacturing value added and 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.

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 life cycle, we study how TFP and aggregate output respond when firms learn faster what they are already learning over their life cycle.

Our methodological approach also differs. Rather than fully specifying a structural model of firm dynamics with parametric stochastic productivity and learning processes, we adopt a formulation inspired by the Arrow-Debreu approach to uncertainty where all possible firm histories are determined at the outset. This minimal structure yields a simple sufficient-statistic characterization of the aggregate effects of accelerating organizational learning across a broad class of models with varying productivity processes, information signals, or ways that firms form beliefs.

AI's Transformative Economic Potential. Current debates about AI's transformative potential (e.g., Agrawal et al., 2025; Chatterji et al., 2025b) have so far focused on automation of production (e.g., Acemoglu, 2025; Ide and Talamàs, 2025; Ide, 2025; Beraja and Zorzi, 2025) — often 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 focus on AI's ability to accelerate organizational learning. A key advantage of this approach is that the statistics required to quantify the aggregate impact are readily available in existing firm-level data. The implied aggregate effects are substantial, suggesting that AI's transformative potential as an organizational learning technology may be large. In addition, because VOLT is largely orthogonal to standard measures of AI exposure (Eloundou

et al., 2024), it complements them in identifying which economic activities are most likely to be transformed by AI.

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. Similarly, Noy and Zhang (2023) and Dell’Acqua et al. (2023) find that the productivity gains from generative AI are concentrated among lower-performing workers. 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.

We complement this literature by measuring AI’s potential for aggregate 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 productivity and output? Our results highlight that the economic impact may arise not only through the within-organization productivity gains emphasized in this literature, but also through lifespan effects that increase the number of active firms in equilibrium. Our analysis further shows that both the productivity and lifespan effects are attenuated by decreasing returns to organizational capital — reflected by the exponent $1 - \nu$ in the VOLT formula (1.1) — resulting in macroeconomic impacts that are an order of magnitude smaller than micro-level estimates would suggest.

2 Model

We describe a general class of models of firm dynamics in the spirit of the canonical framework of Hopenhayn (1992) — where firm heterogeneity, entry, and exit are central — but also incorporate organizational learning. We focus on learning processes that build organizational capital and progressively reveal firm fundamentals. Time is discrete, indexed by $t \in \{0, 1, \dots\}$, and we focus on a stationary equilibrium. Section 5 discusses several extensions of this framework.

2.1 Supply and Demand

Each active firm i produces output y_i with organizational capital z_i and labor n_i :

$$y_i = z_i^{1-\nu} n_i^\nu,$$

where $\nu \in (0, 1)$ denotes the span of control parameter (Lucas, 1978). In Section 3.4, we discuss an alternative interpretation where decreasing returns arise from the demand side. We interpret organizational capital z_i as the collection of internally accumulated, non-tradable assets driving firm productivity (e.g., Prescott and Visscher, 1980; Grant, 1996; Atkeson and Kehoe, 2005; Argote, 2012; Eisefeldt and Papanikolaou, 2013).

As in classical models of firm dynamics (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.² Each firm i observes its organizational capital z_i before choosing n_i , and its static problem is to maximize profits:

$$\pi_i = pz_i^{1-\nu} n_i^\nu - n_i$$

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

$$n^*(z_i) = (pv)^{\frac{1}{1-\nu}} z_i, \quad y^*(z_i) = \frac{1}{pv} n^*(z_i), \quad \pi^*(z_i) = \frac{1-\nu}{\nu} n^*(z_i). \quad (2.1)$$

As a baseline, we take the aggregate demand that firms face to be fixed at a given level of expenditure E :

$$pY = E, \quad (2.2)$$

where Y denotes aggregate output $\int_i y_i di$.

We can interpret this condition as reflecting an aggregate economy where the labor supply is constant. This is because, 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 . Alternatively, as we discuss in Section 4.3, we can interpret the model as that of an industry facing a demand function with unit elasticity, for example as when aggregate output is a Cobb-Douglas aggregator across industries.³

² In Section 5.3, 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.

³ In Section 5.1, we consider an extension in which aggregate demand is given by $p^\theta Y = E$, where θ can capture different elasticities of labor supply (in the aggregate economy) or output demand (in the case of an industry).




Figure 1: The solid line traces the evolution of organizational capital for a given firm type. The cross indicates the age at which the firm chooses to exit the market. The dashed gray lines emanating from each node k illustrate the potential organizational capital paths the firm believes it might follow.

2.2 Firm Types and Information

We adopt a formulation inspired by the Arrow-Debreu approach to uncertainty where we specify the set of all possible firm histories in the economy once and for all at the beginning of time. In particular, we consider different firm types, and we collapse the life cycle of each type into a one time draw of states: a deterministic sequence $\{(z_k, \omega_k)\}_{k=0}^{\infty}$, where (z_k, ω_k) is the firm's state at age k consisting of organizational capital z_k and a probability distribution ω_k over the set \mathcal{S} of all possible types. The firm does not know its type; it learns its state (z_k, ω_k) upon reaching age k , and ω_k captures the firm's beliefs about its type.

Figure 1 illustrates the organizational capital sequence of a specific firm type, alongside the alternative paths of organizational capital it considers possible at each age.

This Arrow-Debreu approach contrasts with the sequential formulations in the literature on firm dynamics, which fully specify productivity and learning processes, and has three main advantages. First, generality: it applies to a wide range of environments with different productivity dynamics, information signals, or ways that firms form beliefs. Second, convenience: it provides a natural formulation to conceptualize changes in organizational learning technologies (Section 3.1). Third, transparency: it highlights the minimum assumptions needed to obtain our sufficient statistic result (3.3).

2.3 Exit and Lifespans

We denote by $V(z_k, \omega_k, p)$ the value of a firm at the beginning of age k given its state (z_k, ω_k) and the output price p . Each period, the firm earns current profits and then chooses whether to continue operating or exit. Firm i 's value function therefore satisfies:

$$V(z_k, \omega_k, p) = \pi^*(z_k, p) + \beta \max(\mathcal{O}(z_k, \omega_k, p), \mathbb{E}_{\omega_k}[V(z_{k+1}, \omega_{k+1}, p)]), \quad (2.3)$$

where $\mathcal{O}(\cdot)$ denotes an outside option, the expectation uses the firm's beliefs ω_k , and $\beta < 1$ is the discount factor. Each firm exits the first time the expected continuation value falls below the outside option $\mathcal{O}(\cdot)$.

We take the outside option $\mathcal{O}(\cdot)$ to be proportional to $p^{\frac{1}{1-\nu}}$. This would arise naturally, for example, if the outside option involved producing the final good.⁴ In this case, both profits and outside options scale with the output price in the same way, so value functions do too and exit decisions are independent of the price. This has two desirable properties. First, changes in aggregate or industry expenditures — for example, coming from population changes or preferences — do not lead to changes in exit rates in the stationary equilibrium. Second, the lifespan of a firm becomes an intrinsic property of its type S . We denote the lifespan of firm type S by $\ell(S)$ and, to guarantee a well-defined stationary distribution of firm types, we assume that firm lifespans are finite: $\ell(S) < \infty$ for all possible types. Also, with some abuse of notation, we denote the type specific component of firm values at entry by $V_0(S) \equiv \frac{V(z_0(S), \omega_0(S), p)}{p^{\frac{1}{1-\nu}}}$.

2.4 Entry

A continuum of firms enters the economy in each period, with the mass of entrants normalized to one.⁵ Potential entrants can partition the set \mathcal{S} of all possible types. Each entrant selects the subset with highest expected value and draws a type from that subset according to a probability distribution μ . For simplicity, we assume that — before the introduction of the organizational learning technology described in Section 3 — entrants can identify only a single subset (namely, the full set \mathcal{S}).

One interpretation of this setup is that the flow of potential entrepreneurs that can

⁴ For example, suppose that the outside option involves the firm producing for a number of periods with some productivity $z_o(S)$. One interpretation is that firms do not exit right away, but they instead wind down their operations slowly.

⁵ As we discuss in Section 6.1, the organizational learning technology we consider does not change the entrants' expected lifetime flow of profits so, for our purposes, assuming a fixed mass of entrants is a reasonable simplifying assumption. In that section, we also discuss an alternative specification in which the number of firms — rather than the flow of entrants — is constant, as in [Atkeson and Kehoe \(2005\)](#).

generate ideas is scarce. Another interpretation is that the flow of financing is scarce: each period there are many entrepreneurs ready to implement different ideas, but there are only funds to start a unit mass of firms, and venture capitalists direct resources toward types in the subset that they can identify with the highest expected value.

2.5 Organizational Learning

In order to conceptualize organizational learning, we introduce an age threshold τ to partition firms into “young” (age $k \leq \tau$) and “mature” (age $k > \tau$). We also partition the set of all types \mathcal{S} into two disjoint sets:⁶

- *Exitors* $\mathcal{S}_{\leq \tau} := \{S \in \mathcal{S} \mid \ell(S) - 1 \leq \tau\}$: The set of types that exit while young.
- *Survivors* $\mathcal{S}_{> \tau} := \{S \in \mathcal{S} \mid \ell(S) - 1 > \tau\}$: The set of types that survive to maturity.

We consider two forms of organizational learning beyond the learning captured by the information states ω_k — both of them consistent with the empirical regularity that mature firms are more productive than young firms (Hsieh and Klenow, 2014). The first one captures learning in the form of organizational capital accumulation. We denote by $z_k(S)$ the organizational capital of firm type S at age k .

Assumption 1. (Organizational Capital Accumulation). *Survivors’ average organizational capital while young is smaller than their average organizational capital upon reaching maturity:*

$$\mathbb{E}_\mu \left[\frac{\sum_{k=0}^{\tau} z_k(S)}{\tau + 1} \mid S \in \mathcal{S}_{> \tau} \right] \leq \mathbb{E}_\mu [z_{\tau+1}(S) \mid S \in \mathcal{S}_{> \tau}]$$

The second assumption is that survivors are positively selected in terms of their expected value upon entry. This assumption captures the idea that early exit disproportionately removes firms with weak organizational capital or poor growth prospects.

Assumption 2. (Survivors Are Positively Selected). *Survivors’ expected value is larger than that of exitors:*

$$\mathbb{E}_\mu [V_0(S) \mid S \in \mathcal{S}_{> \tau}] \geq \mathbb{E}_\mu [V_0(S) \mid S \in \mathcal{S}_{\leq \tau}].$$

3 VOLT: A Sufficient Statistic Result

We now describe a counterfactual economy in which a technology accelerates organizational learning, and derive a sufficient statistic for the resulting value of organizational

⁶ Type S exits at age $\ell(S) - 1$, so it exits at age $k \leq \tau$ if $\ell(S) - 1 \leq \tau$.

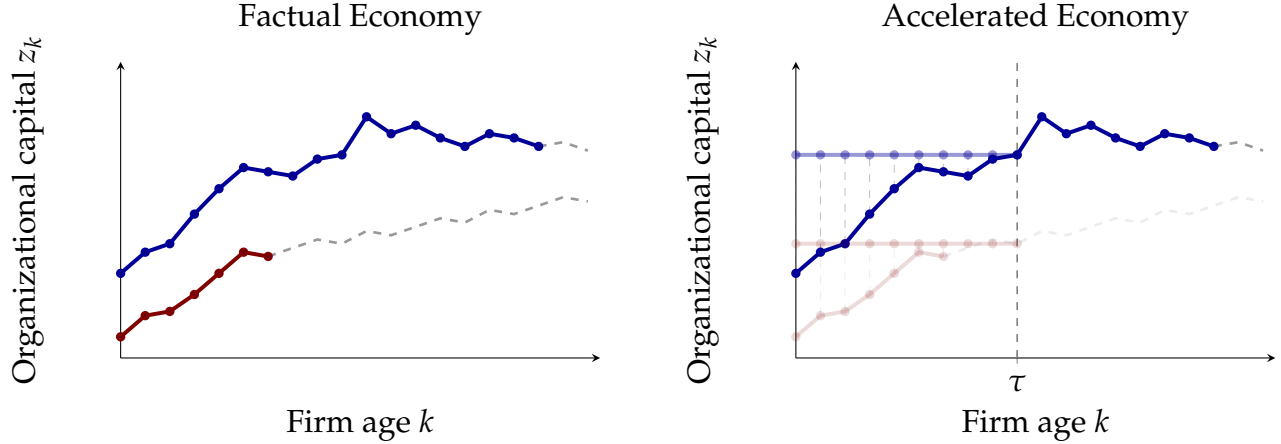


Figure 2: Accelerating early organizational learning in a two-type economy. First, young firms accumulate organizational capital faster, effectively endowing them with the average productivity they would achieve upon reaching maturity. Second, the technology generates early signals that allow entrants to identify survivors, preventing the entry of exitor types.

learning technologies (VOLT). We focus on the long-run aggregate effects once the technology is widely adopted and the economy has converged to a new stationary equilibrium, abstracting from adoption costs and transitional dynamics.

3.1 Accelerating Organizational Learning

We consider a counterfactual economy with faster accumulation of organizational capital and improved prediction about fundamentals before entry. Figure 2 illustrates this counterfactual in a two-type economy. Combining both learning channels is crucial because it preserves the optimal exit decisions of all types that enter the counterfactual economy, allowing us to use observed life-cycle dynamics to evaluate it. In Section 5.4, we discuss the extent to which these two learning channels can be disentangled.

3.1.1 Accelerating organizational capital accumulation

The technology accelerates a firm’s organizational capital accumulation while young. Specifically, the technology changes the mapping from firm types to organizational capital: for each age $k \leq \tau$, the counterfactual organizational capital of type S at age k weakly increases $z_k^c(S) \geq z_k(S)$, subject to the following three constraints.

First, we impose the following aggregate consistency condition: the average counterfactual organizational capital of survivors while young cannot exceed their organizational

capital $E_\mu [z_{\tau+1}(S) \mid S \in \mathcal{S}_{>\tau}]$ upon reaching maturity. To capture the full potential of organizational learning technologies, we assume this bound is reached:⁷

$$\mathbb{E}_\mu \left[\frac{\sum_{k=0}^{\tau} z_k^c(S)}{\tau + 1} \mid S \in \mathcal{S}_{>\tau} \right] = E_\mu [z_{\tau+1}(S) \mid S \in \mathcal{S}_{>\tau}] \quad (3.1)$$

As such, we interpret the acceleration as endowing the average young firm with the organizational capital that would have taken them τ years to accumulate.

Second, to ensure that the acceleration of organizational capital remains independent from the acceleration of learning about fundamentals, we require the counterfactual map $z_k^c(S)$ to not reveal any information that is not already in the firm’s information set at that age, so that beliefs about firm types do not change $\omega_k^c(S) = \omega_k(S)$ for any age k and type S .

Third, we require that the expected value of survivors in the counterfactual continues to dominate that of exitors:

$$\mathbb{E}_\mu [V_0^c(S) \mid S \in \mathcal{S}_{>\tau}] \geq \mathbb{E}_\mu [V_0^c(S) \mid S \in \mathcal{S}_{\leq\tau}]. \quad (3.2)$$

This condition ensures that the counterfactual organizational capital dynamics do not overturn the fact that survivors are positively selected (Assumption 2).

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 into 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

⁷ In Appendix A, we consider a more general counterfactual that allows for a flexible upper bound on young organizational capital and for the possibility that accelerated learning only partially closes the gap to that bound.

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 accelerates learning about fundamentals prior to entry. In the language of our Arrow-Debreu formulation, the technology provides a finer initial information partition. Specifically, it predicts which types are good enough to survive to maturity, allowing potential entrants to partition the set of types \mathcal{S} into exitors $\mathcal{S}_{\leq\tau}$ and survivors $\mathcal{S}_{>\tau}$. Given Equation 3.2, potential entrants find it optimal to avoid the exitors $\mathcal{S}_{\leq\tau}$, so the distribution of entrants in the counterfactual corresponds to the conditional distribution of μ over survivors $\mathcal{S}_{>\tau}$. In all, we interpret this acceleration as endowing firms with information that would have taken them τ years to learn.

AI and learning about fundamentals. 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.⁸

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 allow firms to obtain years of market feedback using synthetic environments and large-scale data.

⁸ 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).

For example, *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).

3.2 Why Not Other Counterfactuals?

Our approach focuses on counterfactuals that can be evaluated directly using observed data on firm dynamics. This contrasts with fully specified and estimated structural models, which allow for a wide range of counterfactuals but require strong assumptions. The advantage of our approach is that the results are robust to many modeling details regarding productivity processes, information signals, and ways in which firms form beliefs; the limitation is that the set of admissible counterfactuals is restricted.

In particular, admissible counterfactuals must preserve the exit decisions of the firm types that enter the economy. This is why we combine the two learning accelerations, and also why the technology preserves the information structure after entry. Altering this information structure — the mapping $\omega_k(S)$ from firm types to beliefs — would change exit decisions, potentially inducing both earlier and later exits in ways that cannot be directly disciplined by observed data. For instance, the mapping would be altered if a technology allowed firms to obtain more precise signals about their fundamentals after entry. We leave such challenges to future work.

That said, within our class of counterfactuals, we can study more flexible ones, including alternative definitions of mature organizational capital and only partial convergence to it. We analyze these counterfactuals in the Appendix A.

Finally, we can disentangle the two learning accelerations and obtain similar sufficient statistic result, but that requires panel data which is often not publicly available. In particular, we can consider τ_{org} and τ_{fund} as two separate thresholds for accelerating organizational capital accumulation and learning about fundamentals, and we only need that $\tau_{org} \leq \tau_{fund}$. Section 5.4 further discusses how the effects of each acceleration could be isolated.

3.3 VOLT Definition and Sufficient Statistic

We define the Value of Organizational Learning Technologies (VOLT) as the increase in aggregate output that results from accelerating the first τ 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_τ^c is the stationary aggregate output in the counterfactual economy with maturity threshold τ .

We define the ratio of mature to average firm lifespan $\ell_\tau := \bar{\ell}_{>\tau}/\bar{\ell}$, where $\bar{\ell}$ and $\bar{\ell}_{>\tau}$ are the average lifespan of all firms and mature firms, respectively. Similarly, we define the ratio of mature to average firm size $y_\tau := \hat{y}_{>\tau}/\bar{y}$, where \bar{y} is the average output of all firms and $\hat{y}_{>\tau}$ is a weighted average of mature firms' output:

$$\hat{y}_{>\tau} = \frac{\tau+1}{\bar{\ell}_{>\tau}} \bar{y}_{\tau+1} + \left(1 - \frac{\tau+1}{\bar{\ell}_{>\tau}}\right) \bar{y}_{>\tau}.$$

In this weighted average, $\bar{y}_{\tau+1}$ denotes the average output of firms of age $\tau+1$, $\bar{y}_{>\tau}$ denotes the average output of mature firms, and the weight $(\tau+1)/\bar{\ell}_{>\tau}$ corresponds to the share of young firms in the counterfactual economy.⁹

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_τ , and the ratio of mature to average firm lifespan, ℓ_τ . In particular,*

$$\text{VOLT}_\tau = (y_\tau \cdot \ell_\tau)^{1-\nu} \tag{3.3}$$

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

A high value of y_τ indicates that mature firms are substantially more productive than young firms. This gap can arise because organizational capital accumulates slowly — so firms spend many 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 ℓ_τ reflects that exit among young firms is substantially higher than among mature firms. This is symptomatic of an economy where entrants struggle to predict which projects are good enough to reach maturity. In such an environment, technolo-

⁹ The first-order condition 2.1 implies $y_\tau = n_\tau = \hat{n}_{>\tau}/\bar{n}$, where \bar{n} and $\hat{n}_{>\tau}$ are defined analogously to \bar{y} and $\hat{y}_{>\tau}$. This equivalence may not hold when marginal products of labor are not equalized across firms, a question we return to in Section 5.3.

gies that accelerate learning about fundamentals have large effects on aggregate output because they reallocate entrants toward projects with substantially longer lifespans.

Proof of Proposition 1. Aggregating the production function (2.1), we obtain:¹⁰

$$Y = \bar{z}^{1-\nu} M^{1-\nu} N^\nu, \quad (3.4)$$

where M is the mass of firms and $\bar{z} := \frac{1}{M} \int_0^M z_i di$ is the average organizational capital. Following Hopenhayn (2014), we think of firms as a production factor, and refer to $\bar{z}^{1-\nu}$ as total factor productivity (TFP).

Given that labor supply N is fixed (see Section 2.1), we obtain:

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

Computing VOLT_τ therefore only requires evaluating the two ratios $\frac{\bar{z}_\tau^c}{\bar{z}}$ and $\frac{M_\tau^c}{M}$. The key property of the counterfactual that allows us to express these ratios as the factual ratios y_τ and ℓ_τ is that the exit choices of the firms that enter the counterfactual remain exactly as in the factual.

Step 1: Invariance of Exit Choices. First, recall from Section 2.3 that exit decisions are independent of the price level p . Hence, it is enough to consider how the changes in organizational capital and beliefs brought about by the technology affect exit decisions keeping the price p fixed.

Second, 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 $\mathcal{S}_{> \tau}$, and Equation 3.2 guarantees that potential entrants find it optimal to draw their types from the set of survivors $\mathcal{S}_{> \tau}$.

Third, note that, for survivors (the only types present in the counterfactual), exit decisions in the counterfactual are exactly the same as in the factual. For ages $k > \tau$, the technology does not affect exit decisions because it does not change the states S_k for any type S . During ages $k \leq \tau$, by definition, survivors never exit in the factual economy. The higher young organizational capital in the counterfactual relative to the factual — together with the fact that the technology preserves the factual information sets ω_k — guarantees that survivors do not have incentives to exit while young in the counterfac-

¹⁰ Aggregating the first order condition 2.1 gives $N = (p\nu)^{\frac{1}{1-\nu}} \bar{z} M$ and $Y = (p\nu)^{\frac{\nu}{1-\nu}} \bar{z} M$. Combining these two conditions gives Equation 3.4.

tual either.

Step 2. Mapping Counterfactual Aggregates to Factual Data. The cross-sectional distribution of firms at any point in time mirrors the dynamic evolution of a single cohort.¹¹ In particular, the stationary mass M of active firms equals the expected lifespan $\bar{\ell}$:

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

and 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:

$$X = \int_0^M x_i di = \mathbb{E}_\mu \left[\sum_{k=0}^{\ell(S)-1} x_k(S) \right].$$

The two variables that determine aggregate output are the mass of firms M and the average organizational capital \bar{z} . Regarding the mass of firms, given that only survivor types are present in the counterfactual, the ratio M_τ^c / M is equal to $\bar{\ell}_{>\tau} / \bar{\ell}$.

Regarding average organizational capital, in the factual, we have that

$$\bar{z} = \frac{1}{M} \int_0^M z_i di = (vp)^{\frac{-v}{1-v}} \frac{Y}{M} = (vp)^{\frac{-v}{1-v}} \bar{y}. \quad (3.5)$$

Similarly, in the counterfactual, we have $\bar{z}_\tau^c = \frac{1}{M_\tau^c} \int_0^{M_\tau^c} z_i^c di$. Using that the cross-sectional distribution of firms at any point in time mirrors the dynamic evolution of a single cohort, Equation 3.1 implies that

$$\bar{z}_\tau^c = \frac{1}{M_\tau^c} \mathbb{E}_{\mu_{>\tau}} \left[\sum_{k=0}^{\ell(S)-1} z_k^c(S) \right] = \frac{1}{\bar{\ell}_{>\tau}} \mathbb{E}_{\mu_{>\tau}} \left[(\tau+1) z_{\tau+1}(S) + \sum_{k=\tau+1}^{\ell(S)} z_k(S) \right].$$

where $\mathbb{E}_{\mu_{>\tau}}[\cdot]$ denotes $\mathbb{E}[\cdot \mid S \in \mathcal{S}_{>\tau}]$. Rearranging terms and using that $\mathbb{E}_{\mu_{>\tau}}[z_{\tau+1}(S)] = \bar{z}_{\tau+1}$, we get that

$$\bar{z}_\tau^c = \frac{\tau+1}{\bar{\ell}_{>\tau}} \bar{z}_{\tau+1} + \frac{\bar{\ell}_{>\tau} - \tau - 1}{\bar{\ell}_{>\tau}} \bar{z}_{>\tau} = (vp)^{\frac{-v}{1-v}} \hat{y}_{>\tau} \quad (3.6)$$

Combining Equations 3.5 and 3.6, we conclude that $\bar{z}_\tau^c / \bar{z} = y_\tau$. \square

The sufficient statistics y_τ and ℓ_τ capture the two channels through which organiza-

¹¹ This result is a direct application of the stationary population (Palm) identity.

tional learning technologies increase aggregate output: by raising firms’ organizational capital while young — making the average firm more productive — and by allowing entrants to select firm types with longer lifespans — increasing the number of firms in stationary equilibrium and thus alleviating diminishing returns. Both channels increase aggregate organizational capital $M\bar{z}$, so the terms y_τ and ℓ_τ enter Equation 3.3 with exponent $1 - \nu$, reflecting the aggregate decreasing returns to organizational capital (see Equation 3.4). The size effect $y_\tau^{1-\nu}$ corresponds to the increase in TFP, while the lifespan effect $\ell_\tau^{1-\nu}$ corresponds to the contribution of the larger number of firms to aggregate output.

3.4 The Mass of Firms: Decreasing Returns vs. Gains from Variety

In the setting above, decreasing returns to scale arise from the supply side. Under this interpretation, an expansion of the mass of firms shrinks the average firm size and alleviates diminishing marginal returns to scale at the firm level. However, as it is well known (e.g., [Hopenhayn, 2014](#)), this framework is isomorphic to a demand-side interpretation based on monopolistic competition and a constant elasticity of substitution between varieties (see Appendix B). In this case, letting $\sigma > 1$ denote the elasticity of substitution between varieties, the decreasing returns parameter $1 - \nu$ maps to $1/(\sigma - 1)$. While the size effect $y_\tau^{\frac{1}{\sigma-1}}$ still captures the TFP gain from organizational learning technologies, in this case the lifespan effect $\ell_\tau^{\frac{1}{\sigma-1}}$ captures the gains from expanding the set of varieties.

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 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_τ with the employment ratio n_τ . This mapping is natural in our baseline framework because $y_\tau = n_\tau$.¹²

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 46.¹³

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} simply by dividing the total employment over the number of establishments. Similarly, we estimate $\bar{n}_{\tau+1}$ by dividing the total employment in the establishments that are in the bin right above age τ over the number of such establishments, and $\bar{n}_{>\tau}$ by dividing the total employment in mature establishments over the number of such establishments.

Estimating the average lifespan $\bar{\ell}$ and average mature lifespan $\bar{\ell}_{>\tau}$ is more involved because 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 δ_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 .¹⁴

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

¹² 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 5.3.

¹³ 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).

¹⁴ 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.

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, $\sum_{t=0}^{\infty} P(t)$.¹⁵ 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 δ_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 D of the Appendix reports our estimates of $\bar{n}_{>\tau}$, n_{τ} , $\bar{\ell}_{>\tau}$ and ℓ_{τ} for all maturity thresholds allowed by the BDS data.

4.2 VOLT Estimates

Figure 3 illustrates the estimated VOLT_{τ} in the U.S. for all values of the maturity threshold τ supported by the BDS. As a benchmark, we consider a span of control parameter $\nu = .75$ which — under the interpretation of the model with differentiated varieties — corresponds to an elasticity of substitution across varieties of $\sigma = 4$.

Our VOLT 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%. 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 ($\tau = 46$),

¹⁵ In line with our model, this assumes that firms that exit at age k have a lifespan of $k + 1$.




Figure 3: VOLT as a function of the maturity threshold τ . 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 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 ν (equivalently, elasticities of substitution σ).

Figure 4 decomposes the sources of VOLT_{τ} , 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 58% of the total increase suggested by VOLT_{46} , while the size effect $y_{46}^{1-\nu}$ alone is 26% of the total increase suggested by VOLT_{46} . The remaining share of the output increase suggested by VOLT_{46} is due to the interaction between the two effects.¹⁶

These estimates highlight the quantitative importance of the lifespan effect $\ell_{\tau}^{1-\nu}$. Standard analyses of aggregate productivity typically focus on the average productivity of active firms, \bar{z} , which in our interpretation corresponds to average organizational capital *per firm*. However, as emphasized by Hopenhayn (2014), in the presence of decreasing returns to scale, aggregate output depends not only on average productivity per firm but

¹⁶ Section D of the Appendix reports these shares for all maturity thresholds τ .




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

also on the total mass of active firms, M . In particular, an economy may display a small 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 Section 6.1 — where we explore an alternative entry process — and Section 6.3 — where we 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_τ with caution. At first glance, the fact that ℓ_τ exceeds y_τ 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 $\ell_\tau > y_\tau$ 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 5.4, 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 3.3) 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 equilibrium expenditure share — and thus also the labor share — of each industry remains constant. 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.¹⁷ In this figure we assume a span-of-control parameter $\nu = .75$, which corresponds to an elasticity of substitution $\sigma = 4$. While this value is a reasonable benchmark at the aggregate level, estimates of σ within specific sectors are often substantially larger, which would imply a larger ν , and therefore lower values of VOLT. At the same time, if industry demand is elastic (elasticity greater than one), this acts as a countervailing force that increases VOLT (see Section 5.1; in particular, the last paragraph of that section provides a reasonable example in which these two forces offset each other). In any case, to facilitate comparisons with our aggregate VOLT estimates, in this figure we maintain the common benchmark of unit elasticity of demand ($pY = E$) and $\nu = .75$.

The value of organizational learning technologies exhibits significant heterogeneity (Avg. $\text{VOLT}_{46} \approx 1.96$, $\sigma \approx 0.39$) across industries. 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 lifespan and size effects. High-VOLT industries generally exhibit a combination of large lifespan and size effects, though the primary driver varies by sector. The potential

¹⁷ Section E 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.




Figure 5: $\log(\text{VOLT}_{46})$ 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, $\nu = 0.75$.

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 size ($y_{46} \approx 9.6$). In contrast, the high value in Apparel Manufacturing ($\text{VOLT}_{46} \approx 2.9$) is driven disproportionately by its exceptional organizational lifespan effect ($\ell_{46} \approx 22$), whereas Performing Arts ($\text{VOLT}_{46} \approx 2.9$) relies on a balanced contribution from both lifespan and size effects.¹⁸

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 of industry values:

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

We estimate industry shares ψ_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.¹⁹ This yields an economy-wide VOLT of 1.92. 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 finding that organizational learning technologies have the potential to substantially increase aggregate output is not driven by a small number of outlier industries.

5 Extensions

We discuss how our VOLT sufficient statistic result extends to the case of arbitrary demand elasticities (Section 5.1), aggregate productivity growth (Section 5.2), and age-dependent labor wedges (Section 5.3).

¹⁸ These industry-level VOLT estimates are naturally sensitive to the elasticity of demand. We return to this issue in Section 5.1.

¹⁹ 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.

5.1 Demand elasticity

In the baseline class of models described in Section 2, aggregate expenditure is fixed at E . We now generalize this condition by assuming the demand condition takes the form:

$$p^\theta Y = E. \quad (5.1)$$

In an aggregate context, θ serves as a reduced-form parameter capturing the combined income and substitution effects on labor supply.²⁰ In an industry context, θ captures the industry's demand elasticity.

The Value of Organizational Learning Technologies retains a sufficient statistic form:²¹

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

In the baseline case with $\theta = 1$, price declines perfectly offset productivity gains to keep expenditure and employment constant, and the exponent simplifies to $1 - \nu$. 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 $y_\tau \ell_\tau$.

When demand is inelastic ($\theta < 1$), the positive shift 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, the exponent $\frac{\theta(1-\nu)}{\theta(1-\nu)+\nu}$ approaches 0. In this case, the price drop completely offsets the productivity gains, leaving aggregate output unchanged.

This extension highlights that our estimates based on $\theta = 1$ are conservative if organizational learning technologies are applied to an aggregate economy with elastic labor supply, or industries with elastic demand. For example, consider an economy with labor supply $N = p^{-\phi} Y^{-\rho}$, Frisch elasticity ϕ is 0.5 (Chetty et al., 2011) and null income effects (i.e., $\rho = 0$). In this economy, $\theta = 1.5$ so VOLT_{46} goes from 2.08 to 2.66.

At the industry level, reasonable estimates of θ lie between 2 and 8 (Simonovska and Waugh, 2014; Boehm et al., 2023), suggesting that the median VOLT_{46} would be between 2.79 and 6.47, rather than the median of 1.90 in our baseline setting. However, as we

²⁰ For example, if labor supply is $N = p^{-\phi} Y^{-\rho}$, where ϕ is the Frisch elasticity and ρ the income elasticity, then, using the aggregate labor demand condition $pY = \nu N$, we obtain Equation 5.1 with $\theta = \frac{1+\phi}{1+\rho}$.

²¹ To see this, start from the aggregate production function $Y = (\bar{z}M)^{1-\nu} N^\nu$. In equilibrium, employment is a fraction ν of revenue, $N = \nu pY$. Substituting this into the production function yields $Y = \bar{z}M (\nu p)^{\frac{\nu}{1-\nu}}$. Using the inverse demand function $p = E^{\frac{1}{\theta}} Y^{-\frac{1}{\theta}}$ we get $Y \propto (\bar{z}M)^{\frac{\theta(1-\nu)}{\theta(1-\nu)+\nu}}$.

discussed in Section 4.3, in the interpretation of the model with differentiated varieties, σ at the industry level can plausibly be much higher than 4, and this would lead to lower values of VOLT. These two forces can approximately offset each other. For instance, if $\sigma = 12$ (implying $\nu = 11/12$), the resulting VOLT estimates are similar to those in Figure 5 if the demand elasticity is around $\theta = 3.7$, a value well within the range of industry-level estimates.

5.2 Aggregate Productivity Growth

Our baseline analysis abstracts from aggregate productivity growth. We now show that while this is innocuous under disembodied growth, abstracting away from embodied growth can underestimate the relevant size effect of organizational learning technologies.

5.2.1 Disembodied Aggregate Productivity Growth

Suppose the production technology of each firm i at time t is

$$y_{i,t} = A_t z_i^{1-\nu} n_{i,t}^\nu. \quad (5.2)$$

where A_t is productivity parameter common across firms. We focus on a stationary equilibrium in terms of the de-trended price $\tilde{p}_t = A_t p_t$ and aggregate output $\tilde{Y}_t = Y_t / A_t$. Optimal labor demand and profits are proportional to $\tilde{p}_t^{\frac{1}{1-\nu}}$, so ensuring the existence of a stationary equilibrium requires that outside options are also proportional to $\tilde{p}_t^{\frac{1}{1-\nu}}$. In this case, each firm's lifespan $\ell(S)$ is as in the baseline economy without growth, and so is the average organizational capital \bar{z} in stationary equilibrium. Aggregating the stationary output of all firms, we obtain:

$$\tilde{Y} = \bar{z}^{1-\nu} M^{1-\nu} N^\nu.$$

The aggregate demand condition $p_t Y_t = E$ continues to imply that labor supply N is fixed. Hence, if we define VOLT_τ as $\tilde{Y}_\tau^c / \tilde{Y}$, the sufficient-statistic formula (3.3) remains valid. In this case, VOLT is interpreted as the proportional level increase in de-trended aggregate output resulting from accelerated organizational learning.

5.2.2 Embodied Aggregate Productivity Growth

Alternatively, aggregate productivity growth may be embodied in successive technological vintages. To capture this idea, suppose that a firm i born at date b uses blueprint $B_b = (1 + g)^b$ throughout its life. In other words, a firm of age k with state $S_k = (z_k, \omega_k)$

at date t has effective organizational capital $\tilde{z}_{k,t} \equiv B_t (1 + g)^{-k} z_k$, so its production function is:

$$y_k = \tilde{z}_{k,t}^{1-\nu} n_{k,t}^\nu.$$

In this case, age-size profiles reflect both organizational learning and vintage effects. To ensure the existence of a Balanced Growth Path (BGP) with stationary firm age distribution, the outside option at time t of firms born at age b must be proportional to $B_b p_t^{\frac{1}{1-\nu}}$. Under this condition, lifespans remain as in the factual economy. Aggregating the production function, we obtain that along a BGP:

$$Y_t = (\bar{z}_t M)^{1-\nu} N^\nu \quad \text{with} \quad \bar{z}_t = (\nu p_t)^{\frac{-\nu}{1-\nu}} \bar{y}_t$$

The existence of vintages means that Equation (3.1) does not pin down the counterfactual output. For this reason, we need to be explicit about the counterfactual average organizational capital at each young age. Suppose that, at every young age $k \leq \tau$, average organizational capital goes up to the average organizational capital upon reaching maturity:

$$\mathbb{E} [z_k^c(S) \mid S \in \mathcal{S}_{>\tau}] = \mathbb{E} [z_{\tau+1}(S) \mid S \in \mathcal{S}_{>\tau}] \quad \text{for each } k \leq \tau. \quad (5.3)$$

In this counterfactual, we have that $Y_{\tau,t}^c = (\bar{z}_{\tau,t}^c M_\tau^c)^{1-\nu} N^\nu$, so

$$\text{VOLT}_\tau^g \equiv \frac{Y_t^c}{Y_t} = \left(\frac{\bar{z}_{\tau,t}^c M_\tau^c}{\bar{z}_t M} \right)^{1-\nu}$$

The ratio $\frac{M_\tau^c}{M}$ continues to be equal to the lifespan statistic ℓ_τ , which is independent of g . To derive the relevant size statistic capturing $\bar{z}_{\tau,t}^c / \bar{z}_t$, start with $\bar{z}_{\tau,t}^c = \frac{1}{M_\tau^c} \int_0^{M_\tau^c} z_{i,t}^c di$. Using Equation 5.3 and that the cross-sectional distribution of firms at any point in time mirrors the dynamic evolution of a single cohort, we obtain:

$$\bar{z}_{\tau,t}^c = \frac{1}{\bar{\ell}_{>\tau}} \mathbb{E}_{\mu_{>\tau}} \left[\sum_{k=0}^{\tau} B_t (1 + g)^{-k} z_{\tau+1}(S) + \sum_{k=\tau+1}^{\ell(S)} B_t (1 + g)^{-k} z_k(S) \right].$$

where $\mathbb{E}_{\mu_{>\tau}}[\cdot]$ denotes $\mathbb{E}[\cdot \mid S \in \mathcal{S}_{>\tau}]$. Rearranging terms, using the first order condition 2.1, and $\mathbb{E}_{\mu_{>\tau}}[z_{\tau+1}(S)] = \bar{z}_{\tau+1}$, we get:

$$\bar{z}_{\tau,t}^c = (\nu p_t)^{\frac{-\nu}{1-\nu}} \left(\frac{\tau + 1 \sum_{k=0}^{\tau} (1 + g)^{\tau+1-k} \bar{y}_{\tau+1,t}}{\bar{\ell}_{>\tau} (\tau + 1)} + \frac{\bar{\ell}_{>\tau} - \tau - 1}{\bar{\ell}_{>\tau}} \frac{Y_{>\tau,t}}{M_{>\tau,t}} \right) \quad (5.4)$$

We conclude that the embodied-growth analog of Proposition 1 is:

$$\text{VOLT}_\tau^g = \left(\ell_\tau \cdot \frac{\hat{y}_{>\tau,t}^g}{\bar{y}_t} \right)^{1-\nu} \quad \text{where } \hat{y}_{>\tau,t}^g \equiv \frac{\tau+1}{\bar{\ell}_{>\tau}} \left(\frac{\sum_{k=0}^{\tau} (1+g)^{\tau+1-k} \bar{y}_{\tau+1,t}}{\tau+1} \right) + \left(1 - \frac{\tau+1}{\bar{\ell}_{>\tau}} \right) \bar{y}_{>\tau,t}$$

Note that for any given estimates of $\bar{y}_{\tau+1,t}$ and $\bar{y}_{>\tau,t}$, the statistic $\hat{y}_{>\tau,t}^g$ is increasing in g , so ignoring embodied growth understates the size effect. In this sense, our baseline VOLT formula is conservative when there is embodied aggregate productivity growth.

5.3 Wedges

As evidenced by the first-order condition 2.1, employment and output in our baseline model are proportional to organizational capital. This relationship allows us to identify the TFP effect of organizational learning technologies $(\bar{z}_\tau^c/\bar{z})^{1-\nu}$ from the statistic y_τ or its employment analog n_τ . This mapping, however, 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 age, our baseline estimates of the TFP effect of organizational learning technologies based on y_τ or n_τ may be biased.

Age-dependent wedges may arise in a wide range of environments. 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 η_k , so the perceived wage of firms of age k is η_k times the market wage (which we have normalized to 1). Hence, $\eta_k > 1$ corresponds to an age-specific tax on labor, while $\eta_k < 1$ corresponds to an age-specific subsidy. Firms' optimal labor demand and output are

η	0.7	0.8	0.9	1.0	1.1	1.2	1.3
$\text{VOLT}_{46,\eta}$	2.75	2.5	2.28	2.08	1.91	1.75	1.61

Table 1: The Value of Organizational Learning Technologies for several values of relative wedges on young firms when the span-of-control parameter is $\nu = 0.75$.

given by the following modification of the first order condition 2.1:

$$n^*(z_k) = \left(\frac{pv}{\eta_k}\right)^{\frac{1}{1-\nu}} z_k \text{ and } y^*(z_k) = \frac{\eta_k}{pv} n^*(z_k). \quad (5.5)$$

In this distorted economy, the observed output and employment ratios, y_τ and n_τ , no longer coincide, and both mix differences in organizational capital and wedges. Given that we don't observe output in the BDS data, we focus on the employment ratio n_τ . Also, given that our main concern is wedges that affect young and mature firms differently, we focus on the case in which $\eta_k = \eta_{\leq\tau}$ for young firms and $\eta_k = \eta_{>\tau}$ for mature firms.

We denote by $\eta \equiv \eta_{\leq\tau}/\eta_{>\tau}$ the wedge on young relative to mature firms. We define $\text{VOLT}_{\tau,\eta}$ as the increase in aggregate stationary output Y_τ^c/Y created by the accelerated learning technology described in Section 3 *while keeping age-dependent wedges fixed*. Aggregating the first order condition 5.5 and using the aggregate demand condition 2.2, we obtain (see Appendix C):

$$\text{VOLT}_{\tau,\eta} = \left(\frac{\bar{\ell}_{>\tau}}{\bar{\ell}}\right)^{1-\nu} \left(\frac{\alpha_{\leq\tau}^c \eta^{-\frac{\nu}{1-\nu}} \bar{n}_{\tau+1} + \alpha_{>\tau}^c \bar{n}_{>\tau}}{\alpha_{\leq\tau} \eta \bar{n}_{\leq\tau} + \alpha_{>\tau} \bar{n}_{>\tau}}\right) \left(\frac{\bar{n}}{\alpha_{\leq\tau}^c \eta^{-\frac{1}{1-\nu}} \bar{n}_{\tau+1} + \alpha_{>\tau}^c \bar{n}_{>\tau}}\right)^\nu$$

where $\alpha_{\leq\tau}$ denotes the factual share of young firms in stationary equilibrium, and $\alpha_{\leq\tau}^c \equiv \frac{\tau+1}{\bar{\ell}_{>\tau}}$ is its counterfactual analog. In particular, if young firms face larger wedges than mature firms ($\eta > 1$), then our baseline VOLT_τ overestimates the size effect. Table 1 illustrates $\text{VOLT}_{46,\eta}$ as a function of relative distortions η .

5.4 Isolating the Value of Faster Selection

Our framework conceptualizes organizational learning technologies as a dual acceleration: Firms accumulate organizational capital faster, and learn about their underlying fundamentals earlier. A natural question is whether their respective contributions to VOLT can be disentangled.

For example, a tempting counterfactual would accelerate organizational capital accu-

mulation without simultaneously accelerating learning about fundamentals. However, without calibrating a specific model of firm dynamics, such a counterfactual is not empirically disciplined, because we don't know how exit decisions of firms that exit while young would change.

However, under a mild assumption on productivity dynamics, we can derive a lower bound on the value of the selection channel in isolation. To do this, we construct a counterfactual in which learning about fundamentals is accelerated as in Section 3.1.2, while the accumulation of organizational capital remains as in the factual economy: Potential entrants can anticipate which firms exit before maturity — so only survivor types enter — but they follow their factual organizational capital 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. The mass of active firms still increases by the factor ℓ_τ . Note also that the share of mature firms is larger in this counterfactual than in the factual economy. Hence, if — in addition to Assumption 1 — we assume that young survivors are, on average, at least as productive as young exitors, then the average productivity in this counterfactual is at least as large as in the factual economy. Hence, the lifespan effect $\ell_\tau^{1-\nu}$ provides a lower bound on the value of faster selection.

This result implies that the large contribution of the lifespan statistic in our VOLT estimates 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 also generate a lifespan effect.

6 Discussion

We discuss the extent to which the two organizational-learning channels underlying VOLT can be separately identified, the importance of understanding the forces that determine firm creation, how our VOLT estimates relate to standard measures of AI exposure, and the implications of the organizational learning technologies for business dynamism.

6.1 What Determines the Number of Firms?

In the baseline class of models of Section 2, there is a constant flow of entrants. This is a reasonable simplifying assumption in the class of models that we consider because en-

entrants' expected lifetime profits are equal to aggregate profits in the cross section $(1 - \nu) pY$ which — given that $pY = E$ — are the same in the factual and counterfactual economies. The constant entry flow implies a tight relationship between the average lifespan of entrants and the number of firms in a stationary equilibrium, leading to the lifespan effect $\ell_\tau^{1-\nu}$.

Alternatively, 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 of firms is fixed, and the managers' wage (a fixed cost of operation) adjusts endogenously to drive the expected present value of entry to zero.

In this alternative environment, the lifespan effect $\ell_\tau^{1-\nu}$ — which relies on an expansion in the number of active firms — is neutralized, so the value of organizational learning technologies is solely driven by the increase in TFP. However, this productivity effect can be amplified by the least productive firms exiting earlier than in the baseline 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 output price. The combination of higher managerial wages and a lower output price 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 our baseline class of models.

This contrast highlights the importance of understanding the forces that determine firm creation. In our baseline environment with a constant flow of entry, organizational learning technologies affect aggregate output both by raising productivity and by increasing the number of active firms through longer expected lifespans. In contrast, when the number of firms is constrained by the supply of managerial talent, for example, these technologies operate primarily through productivity improvements and selection. Understanding whether the relevant bottleneck lies in entrepreneurial activity, managerial talent, or other factors is thus crucial for evaluating and interpreting the value of organizational learning technologies.

6.2 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 its potential value. In this section, we show that AI exposure metrics are largely uncorrelated to our VOLT estimates, suggesting that VOLT can complement them in assessing which industries may be most affected by AI.

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.

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 Emp_s is total employment in industry s .

Figure 6 plots VOLT_{46} against LLM exposure by sector; Appendix E reports the corresponding numerical estimates. The correlation between the two measures is weak (0.1) and statistically insignificant (the p-value is 0.35), suggesting that VOLT and LLM 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 ($\text{VOLT}_{46} \approx 2.65$, LLM Exposure ≈ 0.38) and Securities & Financial In-

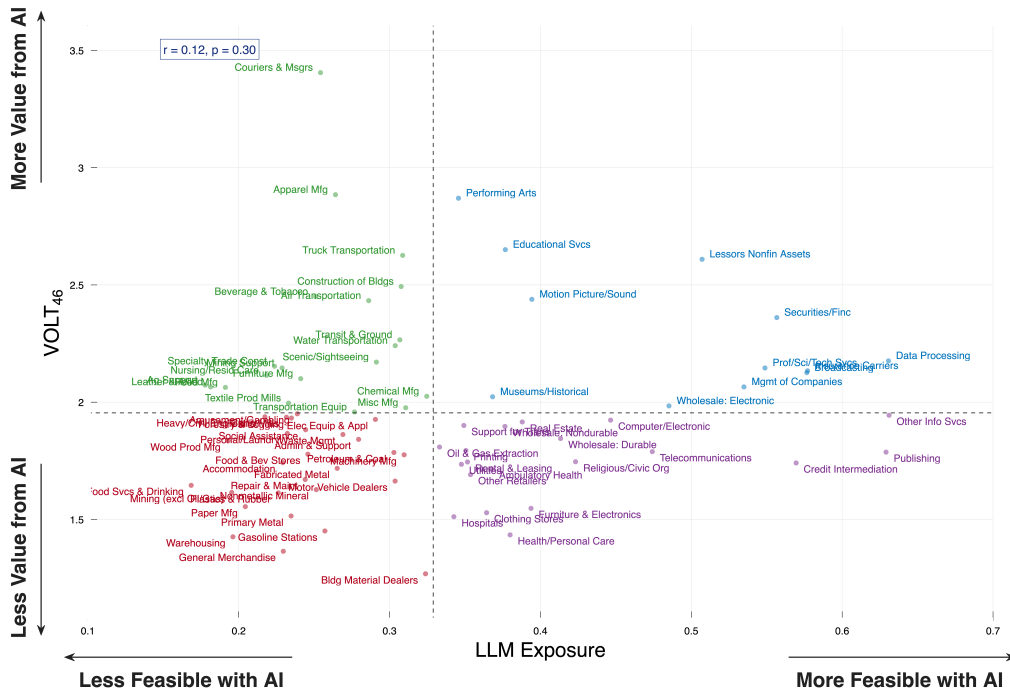


Figure 6: $VOLT_{46}$ vs. LLM Exposure (Eloundou et al., 2024) by 3-Digit NAICS sector. Section E in the Appendix reports the underlying values.

vestments ($VOLT_{46} \approx 2.36$, 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 ($VOLT_{46} \approx 3.41$) but 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.

6.3 Business Dynamism

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 de-

struction, whereby innovation by entrants renders incumbents obsolete and exit reallocates resources toward more productive uses (Aghion and Howitt, 1990). Empirically, this reallocation appears to be 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 than efficient reallocation.

Through this lens, organizational learning technologies reduce 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.

7 Conclusion

Motivated by recent advances in AI, this paper takes a first step toward measuring the value of organizational learning technologies (VOLT). We show that in a broad class of firm dynamics models, VOLT is determined by two simple statistics capturing mature to average firm size and lifespan. This sufficient statistic result provides a transparent way to assess the aggregate output gains from technologies that accelerate organizational learning without relying on a specific model of firm dynamics.

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. Overall, our results point to faster organizational learning as a meaningful and distinct channel of AI’s transformative impact, beyond production automation and scientific discovery.

A general lesson from our analysis is that firm longevity can be an important mechanism mediating the aggregate impact of technologies. In our baseline model, the flow of entrants is fixed, so faster learning raises output not only through higher average productivity but also by extending average firm lifespans — thus increasing the mass of active firms in equilibrium. In contrast, when the number of firms is fixed because of scarce

managerial resources, the longevity channel disappears. This raises a fundamental question: what constrains firm creation? Is it the scarcity of ideas, entrepreneurial energy or financing, implying an inelastic flow of entrants, or is it the scarcity of resources such as managerial ability, implying an inelastic stock of firms? Pinning down the relevant constraint is central for quantifying the gains from faster organizational learning, but more broadly for evaluating other technologies, shocks and policies that affect firm productivity and survival.

Our results also suggest that assessments of the economic impact of a technology — such as AI — should extend beyond the within-organization productivity effects found in microeconomic studies. Our analysis highlights that, to fully capture AI's economic value as an organizational learning technology, it is essential to account for lifespan effects and general equilibrium dampening. In particular, while firm-level treatment effects are informative about mechanisms, our estimates indicate that their aggregate implications can be an order of magnitude smaller once the decreasing returns to organizational capital at the aggregate level are taken into account.

Two limitations of our analysis point to promising directions for future research. First, we have focused on learning processes that build organizational capital and progressively reveal firm fundamentals, abstracting from innovation and creative destruction in the spirit of [Aghion and Howitt \(1990\)](#). Incorporating these forces could reveal how faster organizational learning interacts with the incentives to develop new products (e.g., [Aghion et al., 2017](#); [Trammell and Korinek, 2025](#)).

Second, we have restricted attention to organizational learning technologies that accelerate learning about fundamentals before entry. Accelerating learning about fundamentals after entry would endogenously affect exit decisions in ways that cannot be disciplined with observed factual data without stronger restrictions on the class of models. For example, technologies that improve the precision of post-entry signals could induce both earlier and later exit in the counterfactual. Quantifying the effects of faster post-entry learning about fundamentals is an important challenge for future work.

References

- Daron Acemoglu. The simple macroeconomics of AI. *Economic Policy*, 40(121):13–58, 2025.
- 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.
- 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.

- Isaac Baley and Laura L Veldkamp. *The data economy: Tools and applications*. Princeton University Press, 2025.
- Martin Beraja and Nathan Zorzi. Inefficient automation. *The Review of Economic Studies*, 92(1):69–96, 2025.
- 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, 2025a.
- Ronnie Chatterji, Daniel Rock, and Eduard Talamàs. Transformative AI and the firm. Technical report, National Bureau of Economic Research, 2025b.
- 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.

- Nicolas Crouzet and Janice Eberly. Intangibles, markups, and the measurement of productivity growth. *Journal of Monetary Economics*, 124:S92–S109, 2021.
- 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 Kraymer, 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.
- 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.
- James G March. Exploration and exploitation in organizational learning. *Organization science*, 2(1):71–87, 1991.
- 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.
- 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.

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.

Online Appendix for: The Value of Organizational Learning Technologies

This online appendix contains a detailed description of the quantitative model and its numerical solution, as well as additional results for the article “The Value of Organizational Learning Technologies.”

Any references to equations, figures, or sections that are not preceded by “A.,” “B.,” “C.,” “D.,” or “E.” refer to the main article.

A A More General Accelerated Learning Counterfactual

We have assumed that technologies accelerate learning such that the average organizational capital of young survivors fully reaches their average upon reaching maturity (Equation 3.1). We can broaden this counterfactual to allow for partial accumulation and a more flexible definition of mature organizational capital. Specifically, we can consider a generalized counterfactual such that

$$\mathbb{E}_\mu \left[\frac{\sum_{k=0}^{\tau} z_k^c(S)}{\tau + 1} \mid S \in \mathcal{S}_{>\tau} \right] = \varphi \sum_{k=\tau+1}^{\infty} \omega_k \bar{z}_k \quad (\text{A.1})$$

where $\varphi \leq 1$ governs the extent to which early-stage capital converges to the mature level, $\omega_k \geq 0$ are weights satisfying $\sum_{k=\tau+1}^{\infty} \omega_k = 1$, and \bar{z}_k denotes the cross-sectional average organizational capital of age k firms in the factual economy.

This formulation coincides with our baseline condition when full accumulation occurs ($\varphi = 1$) and the mature benchmark is defined by the initial mature period ($\omega_{\tau+1} = 1$). Under this generalized specification, our sufficient statistic result for VOLT continues to hold, provided we redefine the mature output statistic $\hat{y}_{>\tau}$ as:

$$\hat{y}_{>\tau} = \frac{\tau + 1}{\bar{\ell}_{>\tau}} \left(\varphi \sum_{k=\tau+1}^{\infty} \omega_k \bar{y}_k \right) + \left(1 - \frac{\tau + 1}{\bar{\ell}_{>\tau}} \right) \bar{y}_{>\tau}$$

where \bar{y}_k denotes the cross-sectional average output of age k firms in the factual economy. In particular, a convenient special case is $\omega_k = M_k / \left(\sum_{j=\tau+1}^{\infty} M_j \right)$ and $\varphi = 1$, in which case we get $\hat{y}_{>\tau} = \bar{y}_{>\tau}$.

B Derivation of VOLT under CES Interpretation

Suppose each active firm i produces a differentiated variety q_i with productivity z_i and labor n_i :

$$q_i = a_i n_i$$

The final good is a CES aggregate of varieties:

$$Y = \left(\int_0^M q_i^{\frac{\sigma-1}{\sigma}} di \right)^{\frac{\sigma}{\sigma-1}}, \text{ with } \sigma > 1,$$

while total aggregate expenditure remains fixed at $E = PY$ where P is the CES price index.

Given CES demand, firm i 's revenue is

$$r_i = E \left(\frac{p_i}{P} \right)^{1-\sigma}$$

With marginal cost $1/z_i$, the firm's static pricing problem implies the usual constant markup price

$$p_i = \frac{\sigma}{\sigma - 1} \frac{1}{z_i}.$$

Substituting this pricing rule into demand yields

$$r_i = E \frac{z_i^{\sigma-1}}{\int_0^M z_j^{\sigma-1} dj}.$$

Hence, firm revenue is proportional to $z_i^{\sigma-1}$. Labor demand is proportional to revenue as well, since

$$n_i = \frac{q_i}{z_i} = \frac{\sigma - 1}{\sigma} r_i$$

Therefore, in the CES environment, firm size — measured by revenue or employment — is linear in $z_i^{\sigma-1}$. Using the CES price index,

$$P = \frac{\sigma}{\sigma - 1} \left(\int_0^M z_i^{\sigma-1} \right)^{\frac{1}{1-\sigma}},$$

aggregate output is

$$Y = \frac{E}{\frac{\sigma}{\sigma-1}} (M\bar{z})^{\frac{1}{\sigma-1}}$$

This mirrors the baseline aggregation result $Y \propto (M\bar{z})^{1-\nu}$, except that the exponent $1 - \nu$ is replaced by $\frac{1}{\sigma-1}$.

C Derivation of VOLT in the Presence of Wedges

From the firm's static profit maximization problem, we obtain

$$n_i = \left(\frac{p\nu}{\eta_i} \right)^{\frac{1}{1-\nu}} z_i \quad \text{and} \quad y_i = \left(\frac{p\nu}{\eta_i} \right)^{\frac{\nu}{1-\nu}} z_i$$

Aggregating, we obtain that

$$N = (pv)^{\frac{1}{1-\nu}} M \sum \alpha_k \eta_k^{-\frac{1}{1-\nu}} \bar{z}_k \text{ and } Y = (pv)^{\frac{\nu}{1-\nu}} M \sum \alpha_k \eta_k^{-\frac{\nu}{1-\nu}} \bar{z}_k$$

In particular, $(pv)^{\frac{\nu}{1-\nu}} = N^\nu \left(M \sum \alpha_k \eta_k^{-\frac{1}{1-\nu}} \bar{z}_k \right)^{-\nu}$, so we get

$$Y = M \sum \alpha_k \eta_k^{-\frac{\nu}{1-\nu}} \bar{z}_k N^\nu \left(M \sum \alpha_k \eta_k^{-\frac{1}{1-\nu}} \bar{z}_k \right)^{-\nu} = M^{1-\nu} \frac{\sum \alpha_k \eta_k^{-\frac{\nu}{1-\nu}} \bar{z}_k}{\left(\sum \alpha_k \eta_k^{-\frac{1}{1-\nu}} \bar{z}_k \right)^\nu} N^\nu$$

Given that N is fixed, we obtain:

$$\frac{Y_\tau^c}{Y} = \left(\frac{M_\tau^c}{M} \right)^{1-\nu} \frac{\sum \alpha_k^c \eta_k^{-\frac{\nu}{1-\nu}} \bar{z}_k^c}{\sum \alpha_k \eta_k^{-\frac{\nu}{1-\nu}} \bar{z}_k} \left(\frac{\sum \alpha_k \eta_k^{-\frac{1}{1-\nu}} \bar{z}_k}{\sum \alpha_k^c \eta_k^{-\frac{1}{1-\nu}} \bar{z}_k^c} \right)^\nu \quad (\text{A.1})$$

Young firms (age $k \leq \tau$) face wedge $\eta_{\leq \tau}$ and mature firms (age $k > \tau$) face wedge $\eta_{> \tau}$. To compute Y , we must map the unobservable organizational capital \bar{z}_k to observable average employment \bar{n}_k , using:

$$\bar{z}_{\leq \tau} \propto \eta_{\leq \tau}^{\frac{1}{1-\nu}} \bar{n}_{\leq \tau}, \quad \bar{z}_{> \tau} \propto \eta_{> \tau}^{\frac{1}{1-\nu}} \bar{n}_{> \tau}, \quad \text{and} \quad \bar{z}_{\tau+1} \propto \eta_{> \tau}^{\frac{1}{1-\nu}} \bar{n}_{\tau+1}$$

Defining the relative distortion $\eta \equiv \eta_{\leq \tau} / \eta_{> \tau}$, we obtain:

$$\begin{aligned} \sum \alpha_k \eta_k^{-\frac{\nu}{1-\nu}} \bar{z}_k &= \alpha_{\leq \tau} \eta_{\leq \tau}^{-\frac{\nu}{1-\nu}} \left(\eta_{\leq \tau}^{\frac{1}{1-\nu}} \bar{n}_{\leq \tau} \right) + \alpha_{> \tau} \eta_{> \tau}^{-\frac{\nu}{1-\nu}} \left(\eta_{> \tau}^{\frac{1}{1-\nu}} \bar{n}_{> \tau} \right) \\ &= \alpha_{\leq \tau} \eta_{> \tau}^{\frac{\eta_{\leq \tau}}{\eta_{> \tau}}} \bar{n}_{\leq \tau} + \alpha_{> \tau} \eta_{> \tau} \bar{n}_{> \tau} \\ &= \eta_{> \tau} (\alpha_{\leq \tau} \eta \bar{n}_{\leq \tau} + \alpha_{> \tau} \bar{n}_{> \tau}) \end{aligned}$$

and

$$\begin{aligned} \sum \alpha_k \eta_k^{-\frac{1}{1-\nu}} \bar{z}_k &= \alpha_{\leq \tau} \eta_{\leq \tau}^{-\frac{1}{1-\nu}} \left(\eta_{\leq \tau}^{\frac{1}{1-\nu}} \bar{n}_{\leq \tau} \right) + \alpha_{> \tau} \eta_{> \tau}^{-\frac{1}{1-\nu}} \left(\eta_{> \tau}^{\frac{1}{1-\nu}} \bar{n}_{> \tau} \right) \\ &= \alpha_{\leq \tau} \bar{n}_{\leq \tau} + \alpha_{> \tau} \bar{n}_{> \tau} = \bar{n} \end{aligned}$$

In the counterfactual economy, the learning technology endows young firms with mature organizational capital ($\bar{z}_{\tau+1}$) but leaves the age-dependent distortions unchanged.

The counterfactual aggregates are evaluated using the counterfactual firm shares (α_k^c):

$$\begin{aligned}
\sum \alpha_k^c \eta_k^{-\frac{v}{1-v}} \bar{z}_k^c &= \alpha_{\leq \tau}^c \eta_{\leq \tau}^{-\frac{v}{1-v}} \bar{z}_{\tau+1} + \alpha_{> \tau}^c \eta_{> \tau}^{-\frac{v}{1-v}} \bar{z}_{> \tau} \\
&= \alpha_{\leq \tau}^c \eta_{\leq \tau}^{-\frac{v}{1-v}} \left(\eta_{> \tau}^{\frac{1}{1-v}} \bar{n}_{\tau+1} \right) + \alpha_{> \tau}^c \eta_{> \tau}^{-\frac{v}{1-v}} \left(\eta_{> \tau}^{\frac{1}{1-v}} \bar{n}_{> \tau} \right) \\
&= \alpha_{\leq \tau}^c \eta_{> \tau} \left(\frac{\eta_{\leq \tau}}{\eta_{> \tau}} \right)^{-\frac{v}{1-v}} \bar{n}_{\tau+1} + \alpha_{> \tau}^c \eta_{> \tau} \bar{n}_{> \tau} \\
&= \eta_{> \tau} \left(\alpha_{\leq \tau}^c \eta^{-\frac{v}{1-v}} \bar{n}_{\tau+1} + \alpha_{> \tau}^c \bar{n}_{> \tau} \right)
\end{aligned}$$

Similarly,

$$\begin{aligned}
\sum \alpha_k^c \eta_k^{-\frac{1}{1-v}} \bar{z}_k^c &= \alpha_{\leq \tau}^c \eta_{\leq \tau}^{-\frac{1}{1-v}} \bar{z}_{\tau+1} + \alpha_{> \tau}^c \eta_{> \tau}^{-\frac{1}{1-v}} \bar{z}_{> \tau} \\
&= \alpha_{\leq \tau}^c \eta_{\leq \tau}^{-\frac{1}{1-v}} \left(\eta_{> \tau}^{\frac{1}{1-v}} \bar{n}_{\tau+1} \right) + \alpha_{> \tau}^c \eta_{> \tau}^{-\frac{1}{1-v}} \left(\eta_{> \tau}^{\frac{1}{1-v}} \bar{n}_{> \tau} \right) \\
&= \alpha_{\leq \tau}^c \eta^{-\frac{1}{1-v}} \bar{n}_{\tau+1} + \alpha_{> \tau}^c \bar{n}_{> \tau}
\end{aligned}$$

Because the total mass of active firms is proportional to the expected lifespan, $M_\tau^c / M = \bar{\ell}_{> \tau} / \bar{\ell}$. Plugging the evaluated summations into Equation A.1 gives

$$\frac{Y_\tau^c}{Y} = \left(\frac{\bar{\ell}_{> \tau}}{\bar{\ell}} \right)^{1-v} \left(\frac{\alpha_{\leq \tau}^c \eta^{-\frac{v}{1-v}} \bar{n}_{\tau+1} + \alpha_{> \tau}^c \bar{n}_{> \tau}}{\alpha_{\leq \tau} \eta \bar{n}_{\leq \tau} + \alpha_{> \tau} \bar{n}_{> \tau}} \right) \left(\frac{\bar{n}}{\alpha_{\leq \tau}^c \eta^{-\frac{1}{1-v}} \bar{n}_{\tau+1} + \alpha_{> \tau}^c \bar{n}_{> \tau}} \right)^v$$

D Aggregate Sufficient Statistics and VOLT

Table 2: Aggregate Results by Maturity Threshold; Baseline $\nu = 0.75$.

Threshold (τ)	VOLT	ℓ_τ	y_τ	$\bar{\ell}_{>\tau}$	$\bar{n}_{>\tau}$	$\frac{\ell_\tau^{1-\nu}-1}{VOLT-1}$	$\frac{y_\tau^{1-\nu}-1}{VOLT-1}$
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 3: Aggregate Results by Maturity Threshold; $\nu = .7$

Threshold (τ)	VOLT	ℓ_τ	y_τ	$\frac{\ell_\tau^{1-\nu}-1}{VOLT-1}$	$\frac{y_\tau^{1-\nu}-1}{VOLT-1}$
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 4: Aggregate Results by Maturity Threshold; $\nu = .8$

Threshold (τ)	VOLT	l_τ	y_τ	$\frac{l_\tau^{1-\nu}-1}{VOLT-1}$	$\frac{y_\tau^{1-\nu}-1}{VOLT-1}$
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

E Industry Level VOLT and LLM Exposure

Table 5: Industry-level estimates of VOLT, its components, and LLM exposure.

Code	Industry	VOLT ₄₆	ℓ_{46}	y_{46}	$\ell_{46}^{1-\nu}$	$y_{46}^{1-\nu}$	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	Furniture & Electronics	1.55	5.50	1.04	1.53	1.01	0.39
444	Bldg Material Dealers	1.27	3.87	0.67	1.40	0.90	0.32

Continued on next page

Code	Industry	VOLT ₄₆	ℓ_{46}	y_{46}	$\ell^{1-\nu}$	$y^{1-\nu}$	LLM Exposure
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	Other Retailers	1.69	6.87	1.17	1.62	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 & Msgs	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
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

Continued on next page

Code	Industry	$VOLT_{46}$	ℓ_{46}	y_{46}	$\ell^{1-\nu}$	$y^{1-\nu}$	LLM Exposure
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