

The Value of Codifying Organizational Knowledge*

Martin Beraja

UC Berkeley

Eduard Talamàs

IESE Barcelona

September 9th, 2026

Abstract

Mature organizations are more productive largely because of knowledge accumulated through experience. Most of it remains tacit, so new organizations cannot easily learn it from their predecessors. Artificial intelligence can codify such knowledge without humans articulating it. How much is at stake? We define the value of codifying organizational knowledge as the output gain when new organizations are “born old,” drawing their initial knowledge from the distribution of mature ones. In a class of knowledge-based models of firm dynamics, this value depends on two statistics from observed firm life cycles — the relative size and lifespan of mature firms — and a modest AI system learning from successive cohorts can achieve it. U.S. data suggest a potential gain of up to one GDP, implying that AI could have a meaningful macroeconomic impact by codifying organizational knowledge alone, even if it automated neither production nor innovation. Across industries, the value of codification varies widely and complements existing measures of AI exposure.

1 Introduction

Knowledge that accumulates slowly, over years of operating, explains much of why mature organizations are more productive (Prescott and Visscher, 1980; Jovanovic and Nyarko, 1995; Atkeson and Kehoe, 2005; Eisefeldt and Papanikolaou, 2013). New organizations cannot simply buy or rent that knowledge. They must learn many of the same

*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.

lessons through the same trial and error, retracing much of the same path, often failing before they do. The reason is that most organizational knowledge remains tacit: it cannot be easily articulated or written down anywhere, it lives in the organization's routines and people, and it moves within the organization far more easily than from one organization to another (Polanyi, 1966; Nonaka, 1994; Zander and Kogut, 1995; Szulanski, 1996).

Codification lets organizational knowledge diffuse without moving the people who hold it, so that new organizations can learn from their predecessors (Teece, 1981; Cowan and Foray, 1997). Efforts to codify organizational knowledge have taken many forms: manuals and blueprints record technical knowledge (Mokyr, 2002; Juhász et al., 2025), scientific management breaks work into measured tasks (Taylor, 1911), franchising specifies and replicates routines across establishments (Darr et al., 1995; Argente et al., 2025), and management consulting carries best practices from old to new organizations (Bloom et al., 2016). These codification technologies rely on humans being able to articulate knowledge explicitly: they diffuse what someone can put into words. The rest of organizational knowledge has so far resisted codification (Winter et al., 2012).

Artificial intelligence (AI) technologies codify knowledge without requiring human articulation. They learn directly from examples of practice — operating records, interaction logs, demonstrations, and expert feedback — imitating human experience and generalizing from it (Silver and Sutton, 2025). They can therefore codify, cheaply and at scale, organizational knowledge that until now was too costly or impossible to write down as instructions a human could read (Brynjolfsson et al., 2018; Autor, 2024; Brynjolfsson et al., 2025). The food company NotCo is a concrete example: its AI system accumulated knowledge through years of interaction with the firm's chefs and scientists, and the company now offers what it learned to other food firms through a software platform (Chatterji et al., 2026). Anyone who has asked ChatGPT for practical guidance has done a smaller version of the same thing, obtaining useful knowledge built from other people's experience. Google's bid for the internal records of bankrupt Spirit Airlines shows that AI developers are already paying for organizational data (Perrigo, 2026), and Babina et al. (2026) documents that firms are hiring workers with AI skills in organizational jobs.

What is the value of codifying organizational knowledge? Put differently, if AI lives up to the early evidence, how much is at stake? We formalize the view of AI as a codification technology in a broad class of knowledge-based models of firm dynamics à la Atkeson and Kehoe (2005), and show that the potential value of codifying organizational knowledge can be measured using simple statistics from firm life cycles observed in the past. This is a first step toward making precise, and quantifying, a new channel for AI's macroeconomic impact, one that sits alongside the automation of production (Acemoglu

and Restrepo, 2018; Acemoglu, 2025; Beraja and Zorzi, 2025; Ide and Talamàs, 2025, 2026) and of research and innovation (Aghion et al., 2017; Jones, 2025; Davidson et al., 2026; Jones and Tonetti, 2026).

Specifically, we define the value of codifying organizational knowledge as the answer to a thought experiment. Imagine that the knowledge of mature firms in the economy were codified and diffused, so that new firms were “born old:” they would draw their initial knowledge from the distribution of mature firms rather than from the factual distribution. The value V of codifying organizational knowledge is the ratio of stationary output in this counterfactual to stationary output in the factual economy.¹

We then establish two results, one on measurement and one on implementation. First, the value V is determined by two simple statistics: the ratio y of mature to average firm size, and the ratio ℓ of mature firms’ expected remaining lifespan relative to the expected lifespan at entry. In particular:

$$V = (y \cdot \ell)^{1-\nu} \quad (1.1)$$

where ν is the span of control parameter (Lucas, 1978) that governs decreasing returns to scale.² The ratio y captures the increase in the average organizational knowledge of firms, while the ratio ℓ captures the increase in expected lifespans and, in turn, the mass of firms in stationary equilibrium.³ Both effects are amplified as decreasing returns $1 - \nu$ become stronger.

Second, we show that an AI system can implement the counterfactual gradually, without having to observe mature firms’ records and codify their knowledge all at once. A relatively modest technology that observes each cohort’s operating records during its first year, codifies their knowledge and passes it to the next cohort lets every new cohort start one year ahead of its predecessor, and the economy converges to the same stationary equilibrium as in the thought experiment. This implementation diffuses knowledge across firm cohorts, as in Luttmer (2007), rather than through search and imitation, which

¹ Improvement — the gains from experience with existing techniques and products — is bounded but can be large in levels; invention — new techniques and products — is the source of long-run growth (Jovanovic and Nyarko, 1995). Our counterfactual restricts attention to the former, which is consistent with current AI systems that learn by imitating human experience rather than accumulating experience of their own (Silver and Sutton, 2025).

² 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.

³ In the class of models we study, organizational knowledge can raise a firm’s productivity or the demand it faces — a mature firm’s production processes are as much knowledge as its practices for targeting customers or improving quality. We therefore need not take a stand on whether size differences across firms reflect productivity, as in Atkeson and Kehoe (2005), Hsieh and Klenow (2014) and the literature on learning curves (Bahk and Gort, 1993; Benkard, 2000; Levitt et al., 2013), or they reflect demand and quality, as in Foster et al. (2008) and Eslava et al. (2024). The distinction that matters for our purposes is whether the knowledge that mature firms have and entrants lack can be codified.

diffuse it across the productivity distribution (Jovanovic and Rob, 1989; Lucas and Moll, 2014; Perla and Tonetti, 2014; Buera and Oberfield, 2020). Nor does it affect prediction and learning by incumbents — the mechanism studied in models of the data economy (Baley and Veldkamp, 2025) and of firms with imperfect information (David et al., 2016; Baley and Veldkamp, 2023).

Using 2023 U.S. firm-level data from the Business Dynamics Statistics (BDS), we estimate that mature firms are about 7 times larger than the average firm ($y \approx 7$), and their expected remaining lifespan is about 2.2 times the expected lifespan of entrants ($\ell \approx 2.2$). Interpreting these differences as reflecting knowledge accumulated through experience, and taking $\nu = 0.75$ as our baseline span-of-control parameter therefore gives $V \approx (2.2 \times 7)^{1-0.75} \approx 2$, implying that the stationary output gain from organizational knowledge codification is on the order of one GDP. We view this estimate as an upper bound because our thought experiment closes the entire knowledge gap between mature firms and entrants — but the stakes are large enough that AI could have a meaningful macroeconomic impact through codifying existing organizational knowledge, even if it automated neither production nor research and innovation.

We also use the sufficient statistics formula (1.1) to estimate the value of codification across industries. Figure 1 illustrates these estimates and contrasts them with the standard LLM exposure index of Eloundou et al. (2024), which measures how much of an industry’s work current Large Language Models (LLMs) can reach. The four highlighted industries have nearly identical LLM exposure but sharply different values of codification — illustrating how our measure of value can complement widely used measures of AI capability in predicting which industries are likely to be impacted the most.




Figure 1: Similar LLM exposure, different values of codification.

Notes. Each dot represents a three-digit NAICS industry. The sample requires at least 50 deaths among mature firms. The vertical axis reports the value of codification calculated from 2023 BDS firm size and lifespan profiles using $\nu = 0.75$. The horizontal axis reports the employment-weighted occupational exposure measure of [Eloundou et al. \(2024\)](#).

2 A Class of Knowledge-based Models of Firm Dynamics

We now describe a class of models of firm dynamics à la [Atkeson and Kehoe \(2005\)](#), in which firms' investments in organizational knowledge shape their lifecycle. Importantly, entrants draw their initial knowledge from a distribution that reflects organizational knowledge already made portable by existing institutions. Section 3 introduces a new codification technology that allows entrants to draw their knowledge from the distribution of mature firms instead.

For simplicity, the baseline class of models assumes an exogenous flow of entrants, fixed labor supply, and no aggregate growth. Section 5 relaxes these assumptions.

2.1 Production and organizational knowledge

Time is discrete. There is a unit measure of labor and a unit mass of new firms in every period.⁴ An active firm is characterized by its organizational knowledge $z \in \mathbb{R}_+$. A firm with knowledge $z > 0$ chooses employment $n \geq 0$ and learning investment $x \in [0, 1]$. Its output is

$$q = z(1 - x)n^\nu, \quad \nu \in (0, 1),$$

where ν is the span of control parameter (Lucas, 1978).⁵ The learning investment x is the share of current operating capacity devoted to sustaining and developing organizational knowledge rather than current production; it captures activities such as experimentation and retaining experience in organizational routines.

We interpret organizational knowledge z broadly. Above, z raises physical productivity, and hence the firm's revenue. But a mature firm's production processes are as much knowledge as its practices for targeting customers or improving quality. So, equivalently, z could instead raise the demand the firm faces. We return to this case in Section 4.2 where z shifts demand under monopolistic competition and CES preferences.⁶

Organizational knowledge evolves according to

$$z' \sim F(\cdot \mid z, x).$$

The dependence of F on x allows for both endogenous learning and exit. If F improves with x in the sense of first-order stochastic dominance, a firm that invests more expects to know more next period; as knowledge accumulates, the firm grows, which rationalizes the age-size profile in the data. We assume the state $z = 0$ is absorbing and represents exit. We interpret this absorbing state as a loss of organizational capability severe enough to make continued operation unviable; it can result, for example, from the departure of key personnel, the disruption of routines or coordination networks, the loss of records, or the obsolescence of technologies in which knowledge is embedded. Less severe instances of organizational forgetting are represented by transitions to lower positive values of z .⁷

⁴ Section 5 considers elastic entry and labor supply.

⁵ We use labor as the only rented input for simplicity; the results are unchanged if firms also rent physical capital and combine it with labor in a constant-returns composite.

⁶ Whether size differences across firms reflect productivity — as in Bahk and Gort (1993); Benkard (2000); Atkeson and Kehoe (2005); Hsieh and Klenow (2014) and Levitt et al. (2013) — or they reflect demand and quality — as in Foster et al. (2008) and Eslava et al. (2024) — is therefore immaterial for our purposes. What matters is whether the knowledge that mature firms have is the kind that can be codified.

⁷ The management literature uses the term “organizational forgetting” to describe the depreciation of knowledge retained in an organization's members, routines, records, technologies, and networks (Argote, 2013).

Under this assumption, if $x = 0$ leads to $z' = 0$ with probability 1 in the following period, this investment choice corresponds to an endogenous exit choice.

Take output as the numeraire and let w denote the real wage. Conditional on (z, x, w) , labor demand solves

$$\pi(z, x, w) = \max_{n \geq 0} \{z(1-x)n^\nu - wn\}.$$

Optimal employment is

$$n(z, x, w) = \left(\frac{\nu}{w}\right)^{1/(1-\nu)} [z(1-x)]^{1/(1-\nu)}, \quad (2.1)$$

and profits satisfy $\pi(z, x, w) \propto w^{-\nu/(1-\nu)} [z(1-x)]^{1/(1-\nu)}$.

2.2 The firm's dynamic problem

We focus on a stationary equilibrium. The value of a firm with knowledge z satisfies

$$v(z, w) = \max_{x \in [0,1]} \left\{ \pi(z, x, w) + \beta \int v(z', w) F(dz' \mid z, x) \right\}. \quad (2.2)$$

An optimal investment $x > 0$ sacrifices current output to improve the distribution of future organizational knowledge. The assumption that the learning investment is in terms of output makes the tradeoff between current production and future knowledge independent of the wage: As Lemma 1 highlights, after labor is chosen optimally, a wage change rescales current profits and continuation values by the same factor, leaving optimal investment choices unchanged.

Lemma 1. *In a stationary equilibrium, $v(z, w) = w^{-\nu/(1-\nu)} v(z, 1)$.*

Proof. Because $\pi(z, x, w) = w^{-\nu/(1-\nu)} \pi(z, x, 1)$, multiplying the Bellman equation (2.2) by $w^{\nu/(1-\nu)}$ yields the same equation at $w = 1$. Hence $v(z, w) = w^{-\nu/(1-\nu)} v(z, 1)$. The maximization problem is therefore independent of w . $\square$

Let $x^*(z)$ denote the optimal learning investment of a firm with knowledge z . Equilibrium employment and output are proportional to the size index

$$s(z) = [z(1-x^*(z))]^{1/(1-\nu)}.$$

2.3 The stationary firm distribution

Let m_a be the mass of active firms of age a , and h_a be their knowledge distribution. Our normalization of the entry flow to one implies that the mass of age-zero firms is $m_0 = 1$. We interpret h_0 as the distribution of knowledge available to entrants they acquire experience of their own. It includes whatever has already been codified and diffused — through manuals, courses, suppliers, consultants, and other existing formal or informal knowledge transfer institutions — as well as the tacit knowledge that moves with people, through the prior experience of founders and hired workers.

Let $\rho(z)$ denote the probability that, in equilibrium, a firm with knowledge z stays active in the next period:

$$\rho(z) = F(\mathbb{R}_+ \setminus \{0\} \mid z, x^*(z)).$$

The equilibrium mass of firms of age $a + 1$ equals the equilibrium mass of firms of age a times the fraction of such firms that stay active next period

$$m_{a+1} = m_a \int \rho(z) h_a(dz).$$

The knowledge distribution conditional on age follows a similar recursion:

$$h_{a+1}(B) = \frac{\int F(B \mid z, x^*(z)) h_a(dz)}{\int \rho(z) h_a(dz)}, \quad B \subseteq \mathbb{R}_+ \setminus \{0\}.$$

We assume that a firm's expected lifespan at entry — which given that $m_0 = 1$ equals the equilibrium mass of firms — is finite:

$$\bar{\ell} \equiv M \equiv \sum_{a=0}^{\infty} m_a < \infty. \tag{2.3}$$

Stationarity requires the aggregate flow of entry to equal the aggregate flow of exit. Because the entry flow is normalized to one, the exit flow must also equal one. This condition is satisfied by the cohort masses defined above. To see this, note that the mass exiting from the age- a cohort is $m_{a+1} - m_a$. Therefore, the stationary exit flow is

$$\sum_{a=0}^{\infty} (m_a - m_{a+1}) = m_0 - \lim_{a \rightarrow \infty} m_{a+1} = 1$$

where we have used that $m_a \rightarrow 0$, which is implied by Equation 2.3.

Let $\bar{s}$ denote the average size index among active firms:

$$\bar{s} = \frac{1}{M} \sum_{a=0}^{\infty} m_a \int s(z) h_a(dz).$$

Aggregating firms' labor demand (Equation 2.1) and imposing labor-market clearing gives

$$1 = \left(\frac{\nu}{w} \right)^{1/(1-\nu)} \bar{s} M.$$

Equivalently, the equilibrium wage is

$$w = \nu (\bar{s} M)^{1-\nu}.$$

The first-order condition for labor implies $\nu q = wn$. Because aggregate employment equals one, aggregate output is therefore

$$Y = \frac{w}{\nu} = (\bar{s} M)^{1-\nu}.$$

Intuitively, given the decreasing returns within each firm, aggregate output depends both on the average operating capacity and on the number of firms across which labor can be allocated (Hopenhayn, 2014).

2.4 The mature distribution

The counterfactual in Section 3 allows entrants to draw their knowledge from the distribution of mature firms. For this counterfactual to be well defined, such a distribution must exist. We therefore assume that firms' post-entry dynamics admit a unique invariant distribution and that the average operating capacity and survival of factual cohorts converge to the corresponding mature levels. Economically, the assumption is that learning, forgetting, and selection eventually balance in the cross section.

For any distribution h over active states, define the distribution $\mathcal{T}(h)$ after one period of operation conditional on remaining active:

$$\mathcal{T}(h)(B) = \frac{\int F(B \mid z, x^*(z)) h(dz)}{\int \rho(z) h(dz)}.$$

Assumption 1. *There is a unique “mature distribution” h^m over active states such that*

$$\mathcal{T}(h^m) = h^m.$$

Moreover, as $a \rightarrow \infty$,

$$\begin{aligned}\int s(z)h_a(dz) &\longrightarrow \bar{s}^m \equiv \int s(z)h^m(dz) < \infty, \\ \int \rho(z)h_a(dz) &\longrightarrow \bar{\rho}^m \equiv \int \rho(z)h^m(dz) < 1.\end{aligned}$$

3 The Codification Counterfactual

This section defines the value of codifying organizational knowledge, presents our sufficient statistic result to measure this value, and shows how AI can achieve it by codifying only the first-year operating histories of successive cohorts.

3.1 The value of codifying organizational knowledge

We define the value of codifying organizational knowledge V by considering the following thought experiment. A technology codifies what the typical mature firms know and diffuses it to entrants, so that a firm's initial knowledge is drawn from the mature-firm distribution h^m instead of the factual entry distribution h_0 . That is, the counterfactual entry distribution of knowledge is

$$h_0^c = h^m.$$

Figure 2 illustrates the thought experiment. We have in mind a technology that learns from records of organizational practice and makes the resulting capability usable before a new firm develops an operating history of its own. New firms then learn from their predecessors rather than relearning the same lessons through their own investments and trial and error, and they inherit the results of the selection that produced the mature distribution without having to survive that selection themselves.

The value V is the ratio stationary output in this counterfactual to that in the factual economy

$$V = \frac{Y^c}{Y}.$$

The thought experiment is deliberately blunt in one respect and deliberately controlled in another. It is blunt because it closes the entire knowledge gap between entrants and mature firms: entrants come to know as much as mature firms know, though not more, since the dispersion among mature firms is preserved. A more ambitious counterfactual could allow entrants to draw from the knowledge distribution of a selected group




Figure 2: The Thought Experiment

of mature firms; evaluating this counterfactual would require richer data on the selected firms and additional assumptions about how their knowledge evolves.

The experiment is deliberately controlled because the codification technology changes the knowledge available at entry and nothing else. Conditional on its initial knowledge, a firm faces the same production technology and the same process of learning, forgetting, and exit as in the factual economy. Codification makes an empirically observed knowledge distribution available earlier; it neither changes how knowledge is turned into output nor how knowledge evolves after entry, so it delivers a level effect on output rather than a growth effect. This restriction is what lets existing firm life-cycle data discipline the exercise: those data reveal size and exit rates by age, but not how firms would behave under unobserved changes to production or to post-entry knowledge dynamics. Quantifying such changes would require additional evidence and assumptions.

3.2 A sufficient statistics result

Define mature firms' relative size and relative lifespan as

$$y = \frac{\bar{s}^m}{\bar{s}}, \quad \ell = \frac{\bar{\ell}^m}{\bar{\ell}}.$$

Both are observable in the factual economy. Because stationary-equilibrium employment and output are proportional to the size index $s(z)$, y is average employment or output among mature firms relative to the corresponding average among all active firms; ℓ is mature firms' expected remaining lifespan relative to the expected lifespan of an entrant.

Our first result is that these two statistics determine the value of codifying organizational knowledge.

Proposition 1. *The value of codifying organizational knowledge is*

$$V \equiv \frac{Y^c}{Y} = (y \cdot \ell)^{1-\nu}. \quad (3.1)$$

The size term $y^{1-\nu}$ measures the gain from giving all firms the organizational knowledge of mature firms. The lifespan term $\ell^{1-\nu}$ measures the gain from the larger stationary mass of firms generated by giving all firms the mature survival rates.

The proposition rests on two steps. The first is that the counterfactual stationary equilibrium has a simple structure: every cohort is born mature. Entrants draw from $h_0^c = h^m$, and, because $\mathcal{T}(h^m) = h^m$, the knowledge distribution among the surviving members of each cohort remains h^m at every age. Consequently, average size and exit rates are constant across firm age, and aggregation is immediate. The mass of firms of age a is $(\bar{\rho}^m)^a$ so the total stationary mass of firms is $\bar{\ell}^m = (1 - \bar{\rho}^m)^{-1}$, and the average size index at every age is $\bar{s}^m$, so counterfactual output is

$$Y^c = (\bar{\ell}^m \bar{s}^m)^{1-\nu}.$$

The second step is what justifies using the factual $\mathcal{T}(h^m)$ and h^m in the counterfactual. Although codification directly changes the knowledge available at entry, its effects are general equilibrium: more knowledgeable entrants demand more labor and bid up the wage, and a higher wage could in principle alter investment and exit decisions, and with them the survivor operator. Lemma 1 rules this out: in a stationary equilibrium, firms' investment policies do not depend on the wage, so the mapping from a cohort's knowledge distribution to the distribution among its survivors is the same $\mathcal{T}$ in both economies, and h^m remains its invariant distribution. This is what allows the size and survival of mature firms observed in the factual economy to characterize the counterfactual one.

Proposition 1 is silent about the speed of the transition from the factual to the counterfactual stationary equilibrium. Existing firm-life-cycle data discipline the stationary gain because they reveal the relative organizational knowledge and remaining lifespan of mature firms, but not how firms respond along the equilibrium wage path. Disciplining these transition dynamics would require additional evidence or a more tightly specified quantitative model. Applications of Proposition 1 should therefore be interpreted as measuring the potential eventual output gain from codification; how much or how quickly that gain can be realized is a separate question.⁸

⁸ Appendix B makes this distinction concrete. In a fully specified AR(1) example, it derives the adjustment paths and shows how knowledge persistence governs the speed of convergence.




Figure 3: Sequential Implementation

3.3 A sequential implementation with an AI codification technology

We now ask how an AI system could implement the counterfactual and achieve the value in Proposition 1. Read literally, implementing $h_0^c = h^m$ seems to demand a great deal of a codification technology: access to the records of mature firms, which are proprietary, and the capacity to codify decades of accumulated experience at once. Neither is needed, it turns out.

Consider a far more modest AI system that proceeds sequentially. The AI observes only the records of firms in their first year of operation. It codifies their organizational knowledge and makes it available to the next cohort of entrants. That cohort begins with a larger stock of codified knowledge and generates a new set of first-year records from which the technology can learn. Repeating this process can accumulate experience across cohorts even if the technology never observes any firm beyond its first year of operation. Figure 3 illustrates this sequential implementation.

Formally, let H_t denote the entry distribution of the cohort entering at date t , and let $\mathcal{T}_t$ denote the survivor operator: the mapping from a cohort's knowledge distribution at entry to the distribution among its survivors one period later, given the date- t wage and policies. At the end of the cohort's first year, the technology learns from its operating records and makes what it has learned available to the next cohort, so that $H_{t+1} = \mathcal{T}_t(H_t)$. The next proposition shows that this sequential implementation delivers the counterfactual stationary equilibrium in the previous section, and therefore achieves the value of codifying organizational knowledge V in Proposition 1.

Proposition 2. *Suppose an AI codification technology sets $H_{t+1} = \mathcal{T}_t(H_t)$. In a stationary equilibrium, the entry knowledge distribution is h^m , and output equals Y^c .*

Two observations prove the result. First, in any stationary equilibrium with the AI system in place, the wage and firms' policies are constant, so the survivor operator is the

factual $\mathcal{T}$ by Lemma 1, and the entry distribution must satisfy $H = \mathcal{T}(H)$. Second, by Assumption 1, $\mathcal{T}$ has a unique invariant distribution, so $H = h^m$: the only stationary equilibrium the sequential technology admits is the one in which entrants draw from the factual mature distribution.

The proposition does not by itself establish that the economy converges to the stationary equilibrium. Along the transition wages rise, Lemma 1 does not apply, and the operator $\mathcal{T}_t$ depends on date- t policies. The sequence of entry distributions is therefore not the sequence of factual age cross sections $h_0, \mathcal{T}(h_0), \mathcal{T}^2(h_0), \dots$, but a perturbation of it in which each rung is the age- a cross section of the economy with the AI technology in place. Convergence requires that $\mathcal{T}_t \rightarrow \mathcal{T}$ as the wage settles, which we have not shown.

Appendix A proves that convergence nonetheless follows from a simple bounding argument. One economic condition suffices: that the aggregate size index S_t rises along the transition, which holds when the improvement in each cohort's knowledge at entry outweighs the drag of anticipated wage growth on its investment. In Appendix B, we present a fully specified AR(1) example where the economy converges.

4 Estimates

This section takes our sufficient statistic result (Proposition 1) to U.S. data.

4.1 Data, measurement and identification

Data and construction. We use the 2023 release of the Business Dynamics Statistics (BDS), which is constructed from the Longitudinal Business Database and covers U.S. employer firms and establishments (U.S. Census Bureau, 2025). The BDS reports counts, employment, and deaths by age. Our baseline unit is the firm.⁹

In the model, employment is proportional to operating capacity, so we measure y as average employment among mature firms relative to average employment among all active firms. We take as mature the left-censored group born before 1977, whose members are older than 46 years in 2023. To construct lifespans, we follow a synthetic cohort through the BDS age bins, holding the annual death rate constant within each bin.¹⁰ Summing the cohort's survival probabilities gives its expected lifespan, $\bar{\ell}$. For the open-ended

⁹ In Appendix E.1, we treat the establishment as the organizational unit instead.

¹⁰ Deaths reported at age a occurred among firms that would have reached age a had they survived. Thus, if N_a firms survive to age a and D_a firms die before reaching it, the death rate over the preceding interval is $\delta_a = \frac{D_a}{N_a + D_a}$.

mature group, let δ_m denote the average annual death rate among mature firms. We assume that this rate applies at every subsequent age, giving an expected remaining lifespan of $\bar{\ell}^m = 1/\delta_m$.

Identification. We highlight two main identifying assumptions. First, the left-censored group of firms born before 1977 represents the mature distribution, and differences across age groups in the 2023 cross section trace the life cycle of a representative cohort in a stationary environment. Persistent cohort differences or age-specific effects of conditions in 2023 would cause the cross-sectional profile to differ from the underlying life-cycle profile.

Second, we interpret the age profiles of employment and survival as tracing the evolution of organizational knowledge. We use the term in a deliberately broad, functional sense: it encompasses everything an organization can do.¹¹ This interpretation is confounded to the extent that age profiles instead reflect selection on fixed initial differences or changes in the external environment — such as demand, regulation, or input-market conditions — rather than the evolution of organizational capabilities.

Calibration. We take $\nu = 0.75$ as our baseline and report sensitivity to alternative values. [Hopenhayn \(2014\)](#) reviews commonly used calibrations of the parameter ν , noting that span-of-control models often set this parameter to 0.85, while the monopolistic-competition calibration of [Hsieh and Klenow \(2009\)](#) implies $\nu = 2/3$.

4.2 Aggregate estimate

The expected lifespan of a U.S. entrant is 10.4 years. Firms in the mature group have an expected remaining lifespan of 23.4 years, giving $\ell \approx 2.24$. Mature firms employ 174.3 workers on average, compared with 24.6 among all active firms, giving $y \approx 7.09$. Equation 3.1 therefore gives

$$V \approx (7.09 \times 2.24)^{0.25} \approx 2.$$

Thus, our baseline estimates suggest that codification approximately doubles aggregate output. If we take ν to be $2/3$ or 0.85 instead, the value of codification is 2.5 and 1.5, respectively. Thus, across these calibrations, codification raises stationary output by roughly between 50 and 150 percent.

An alternative microfoundation places the relevant curvature in demand rather than

¹¹ Following the tradition of the knowledge-based view of the firm, we set aside the abstract question of what knowledge is — which has intrigued some of the world’s greatest thinkers from Plato to Popper without the emergence of a clear consensus ([Grant, 1996](#)).

production. Suppose each firm produces a differentiated variety with a linear technology, and final output is a CES aggregate of these varieties with elasticity of substitution σ . Under this formulation, the value of codification is $V_\sigma = (y \cdot \ell)^{1/(\sigma-1)}$ (see Appendix C). The estimate $y\ell \approx 16$ therefore implies $V_\sigma \approx 4, 2$, and 1.5 for $\sigma = 3, 5$, and 8 , respectively.¹²

4.3 Industry heterogeneity and AI exposure

We also estimate the value of codification across three-digit NAICS industries. To avoid unstable estimates of mature remaining lifespan, we require complete employment and death data and at least 50 mature-firm deaths in 2023.¹³ This screen leaves 40 industries. Their values of codification have a mean of 1.82, a median of 1.79, and a standard deviation of 0.26, with a range from 1.43 to 2.57.¹⁴

Industry estimates require an additional qualification. The BDS classifies establishments, rather than entire firms, by industry. A firm operating in several industries therefore contributes to each relevant industry, with the same parent-firm age assigned to its establishments. Industry cells consequently overlap: each estimate describes the age profile of activity in that industry, but the estimates cannot be summed to recover the aggregate value.

Table 2 in the Appendix reports the value-of-codification estimates for every industry in our sample. To compare these estimates with the technological reach of current AI, we use the LLM-exposure measure of Eloundou et al. (2024). The authors classify occupational tasks according to whether access to an LLM or LLM-powered software could reduce the time required to complete the task by at least 50 percent while maintaining output quality. We use their intermediate measure, which fully counts tasks directly exposed to an LLM and gives half weight to tasks requiring complementary software, and aggregate it to three-digit NAICS industries using their occupational employment shares. The resulting index measures the share of an industry’s work that current LLM capabilities

¹² The appropriate value of σ depends on the scope of the exercise. Within an industry, it is the elasticity of substitution among firms’ varieties in that industry. At the aggregate level, the relevant curvature must also reflect the more limited substitution across industries, so a within-industry demand elasticity cannot simply be carried over to the economy as a whole.

¹³ We require at least 50 mature-firm deaths because the reliability of a death-rate estimate depends primarily on the number of observed deaths D . Under a binomial benchmark, the relative standard error of both the death rate and its inverse is approximately bounded by $1/\sqrt{D}$. With 50 deaths, this bound is about 14 percent.

¹⁴ Without the death-count screen, the 72 industries with complete disclosed cells have a mean of 1.92, a median of 1.81, a standard deviation of 0.38, and a range from 1.43 to 3.38. The screen removes, among others, several apparent outliers whose mature remaining-lifespan estimates are based on single-digit deaths. Results are similar at nearby cutoffs: thresholds of 25 and 100 deaths retain 51 and 32 industries, respectively, and yield mean values of 1.85 and 1.83 and median values of 1.81 and 1.79.

could substantially accelerate.

Section 1 uses Figure 1 to emphasize that industries with similar LLM exposure can have substantially different values of codification, while industries with similar codification values can differ greatly in their exposure. Figure 4 develops this comparison across the four regions defined by the sample medians, illustrating how the technological reach of current LLMs and the economic value of making mature firms' knowledge available at entry vary independently across industries.




Figure 4: The value of codification against LLM Exposure across industries.

Notes: Each point represents a three-digit NAICS industry. The sample requires complete employment and death data and at least 50 mature-firm deaths in 2023. Dashed lines mark the sample medians.

Industries in the upper-left quadrant exhibit steep size and/or survival profiles — indicating, under our interpretation, that firms acquire substantial productive knowledge through experience that is unavailable to entrants — but relatively little of their work is exposed to current LLMs. Truck transportation, food manufacturing, and building construction generate experience about routing, process control, quality adjustment, project sequencing, and subcontractor performance, while many of the underlying activities remain physical and context dependent.

Industries in the upper-right quadrant combine the same large maturity gap with high LLM exposure. In education, professional services, and wholesale intermediation, student trajectories, completed projects, expert evaluations, and histories of buyer-supplier

matches provide data from which AI can learn, while the life-cycle profiles indicate substantial scope for improving the knowledge available at entry.

Industries below the horizontal median have flatter size and/or survival profiles. Under our interpretation, less productive organizational knowledge is acquired after entry, either because less such knowledge is required or because more of it is already available to entrants. Ambulatory health care, publishing, and printing lie in the lower-right quadrant: many of their tasks are within reach of LLMs, but professional training, protocols, technical standards, and experienced workers already transmit a substantial part of the relevant knowledge across organizations. Specialty trades, repair services, and fabricated-metal manufacturing lie in the lower-left quadrant. Current LLMs have limited reach into their largely physical work, while entrants can obtain much of the requisite expertise by hiring experienced tradespeople, mechanics, and machinists or by using established drawings, codes, and equipment.

5 Extensions

We discuss three departures from the baseline model: endogenous entry, aggregate productivity growth, and elastic labor supply.

5.1 Entry

The baseline normalizes the flow of entrants to one. This keeps the entry margin separate from the effect of codification on the size and survival of each entrant cohort. If the stationary entry flow changes, its contribution can be added directly to the sufficient statistic. Let λ and λ^c denote the factual and counterfactual entry flows. The age distributions defined in Section 2.3 are normalized per entrant, so aggregate output is $Y = (\lambda M \bar{s})^{1-\nu}$. Consequently,

$$\frac{Y^c}{Y} = \left(\frac{\lambda^c}{\lambda} \right)^{1-\nu} (y\ell)^{1-\nu}. \quad (5.1)$$

The term $(\lambda^c/\lambda)^{1-\nu}$ is the additional output effect generated by endogenous entry. A different entry flow changes the equilibrium wage but, by Lemma 1, does not change stationary investment and exit policies.

The direction of the entry response is not determined by y and ℓ alone. Codification raises the capabilities available at entry and therefore raises the private value of entry at a given wage. The additional operating capacity also raises labor demand and the

equilibrium wage, reducing profits at every organizational state. Moreover, codification both brings productive capacity forward and can extend firms' lives, placing additional capacity in more distant periods. These forces can move entry in opposite directions.

One benchmark makes this ambiguity precise. Let $\bar{s}_a = \int s(z)h_a(dz)$ and define the undiscounted and discounted operating capacity generated by one entrant as

$$U = \sum_{a=0}^{\infty} m_a \bar{s}_a, \quad D = \sum_{a=0}^{\infty} \beta^a m_a \bar{s}_a. \quad (5.2)$$

Suppose potential entrants are available elastically and enter until the expected value of entry equals a fixed outside option χ measured in units of labor. With entry flow λ , labor-market clearing is

$$1 = \lambda \left(\frac{\nu}{w} \right)^{1/(1-\nu)} U.$$

The expected value of an entrant is

$$\psi = (1 - \nu) \nu^{\nu/(1-\nu)} w^{-\nu/(1-\nu)} D.$$

Dividing by the wage and using labor-market clearing gives

$$\frac{\psi}{w} = \frac{1 - \nu}{\nu \lambda} \frac{D}{U}. \quad (5.3)$$

The free-entry condition $\psi/w = \chi$ therefore implies

$$\lambda = \frac{1 - \nu}{\nu \chi} \frac{D}{U}. \quad (5.4)$$

The ratio D/U is a capacity-weighted average of β^a . Proportional changes in capacity at every age are absorbed by the wage and leave entry unchanged. Entry responds only to when capacity is generated. Bringing capacity forward raises D/U ; extending the life cycle places more capacity at distant ages and lowers it.

In the complete-codification economy, $m_a^c = (\bar{\rho}^m)^a$ and $\bar{s}_a^c = \bar{s}^m$ at every age. Hence,

$$\frac{D^c}{U^c} = \frac{1 - \bar{\rho}^m}{1 - \beta \bar{\rho}^m},$$

and an unchanged outside option gives

$$\frac{\lambda^c}{\lambda} = \frac{1 - \bar{\rho}^m}{1 - \beta \bar{\rho}^m} \frac{\sum_{a=0}^{\infty} m_a \bar{s}_a}{\sum_{a=0}^{\infty} \beta^a m_a \bar{s}_a}. \quad (5.5)$$

Combining equations (5.1) and (5.5) yields

$$\frac{Y^c}{Y} = (y\ell)^{1-\nu} \left[\frac{1 - \bar{\rho}^m}{1 - \beta \bar{\rho}^m} \frac{\sum_{a=0}^{\infty} m_a \bar{s}_a}{\sum_{a=0}^{\infty} \beta^a m_a \bar{s}_a} \right]^{1-\nu}.$$

This correction uses the full age profile and a discount factor, rather than the two sufficient statistics alone.

As $\beta \rightarrow 1$, both D/U and D^c/U^c converge to one. It follows from equation (5.5) that $\lambda^c/\lambda \rightarrow 1$. In this patient limit, the higher value created by more capable entrants and the wage response exactly offset in the entry condition. The fixed-entry sufficient statistic is therefore also the limiting free-entry calculation. Away from this limit, equation (5.5) provides one explicit entry correction. We retain the fixed entry flow in the baseline because other entry technologies, project-supply elasticities, and startup-resource costs imply different corrections.

5.2 Aggregate productivity growth

For simplicity, our baseline model assumes there is no aggregate growth. We now describe how our sufficient statistic result has or does not have to be adjusted depending on the source of economic growth.

5.2.1 Productivity growth common to all firms

When productivity growth affects all firms in the same proportion, it does not contaminate the age profiles used by the sufficient statistic. It changes the level and growth rate of output, and may change the policies already reflected in the factual life cycle, but it does not change the measured value of codification.

Suppose the production function at date t is

$$q_t = A_t z (1 - x) n^\nu, \quad A_{t+1} = G A_t, \quad G \geq 1, \quad \beta G < 1. \quad (5.6)$$

The factor A_t multiplies the productivity of every firm at date t , regardless of its age. The equilibrium wage grows in proportion to A_t . The ratio A_t/w_t that governs employment is therefore constant over time, while operating profits grow at rate G . Dividing wages and firm values by A_t leaves the same decision problem at every date, with an effective discount factor βG . The normalized Bellman equation is therefore independent of calendar time. Together with the existence and finite-lifespan conditions imposed in Section 2,

this delivers stationary learning and exit policies and a stationary age distribution in detrended units.

Growth can affect firms' learning and exit policies through βG , but the factual life cycles already reflect those policies. The normalized wage remains common to all firms and scales current and future payoffs together, so the invariance argument in Lemma 1 still applies. Because codification leaves G and the other primitives fixed, it does not change the normalized learning or exit policies.

At any date, the same A_t multiplies the operating capacity of every firm. It therefore cancels from the mature-to-average size ratio and from the counterfactual-to-factual output ratio:

$$\frac{Y_t^c}{Y_t} = \frac{Y_t^c / A_t}{Y_t / A_t} = (y \cdot \ell)^{1-\nu}. \quad (5.7)$$

5.2.2 Productivity growth embodied in new vintages

When productivity improvements are embodied in new firms, age is mechanically related to technological vintage: older firms may be less productive even if they have accumulated more organizational knowledge. The raw age profile therefore no longer isolates knowledge accumulation. We show that when the vintage effect is known and multiplicatively separable, it can be removed using the observed age profile and the vintage growth rate.

Suppose that a firm born at date b operates a fixed technological vintage $B_b = G^b$. At date t , a firm of age $a = t - b$ has production function

$$q_{a,t} = B_t G^{-a} z (1 - x) n^\nu. \quad (5.8)$$

Age now measures both organizational experience and distance from the current technological frontier. Conditional on a common wage, employment is proportional to $G^{-a/(1-\nu)} [z (1 - x)]^{1/(1-\nu)}$. Consequently, the raw age-size profile mixes two forces: older firms may have accumulated more organizational knowledge, but they also operate an older vintage. The baseline two-statistic formula cannot be applied to that profile without a correction.

A known multiplicative vintage effect can be corrected because the structure above is separable. The term G^{-a} changes the return to every knowledge state in the same proportion and does not alter the knowledge transition or maintenance technology. After optimizing labor demand, an age- a firm's entire profit stream contains the common factor $B_t G^{-a/(1-\nu)}$. From one period to the next, this factor changes by the same amount for every age. Dividing firm values by it therefore leaves one normalized decision problem, with effective discount factor $\beta G^{-\nu/(1-\nu)}$.

The investment and exit policy can depend on G ; crucially, it does not depend separately on a firm's age. This age invariance is a consequence of the separable vintage specification, not an additional assumption that growth has no effect on firm behavior. The correction is then conceptually simple: use the known vintage growth rate to separate the vintage component from each point of the observed age profile, and apply that component to the counterfactual age distribution. Let $\bar{s}_a \equiv \int s(z)h_a(dz)$ and define the growth-adjusted factual capacity stock

$$S_g \equiv \sum_{a=0}^{\infty} m_a G^{-a/(1-\nu)} \bar{s}_a.$$

Under complete codification, the corresponding stock is

$$S_g^c = \bar{s}^m \sum_{a=0}^{\infty} (\bar{\rho}^m G^{-1/(1-\nu)})^a = \frac{\bar{s}^m}{1 - \bar{\rho}^m G^{-1/(1-\nu)}}.$$

Hence the value of codification for detrended output is

$$V_g = \left(\frac{S_g^c}{S_g} \right)^{1-\nu}. \quad (5.9)$$

Appendix D.1 derives these expressions. Unlike the baseline sufficient statistic, the correction uses the full age profile and the vintage growth rate rather than the two raw moments y and ℓ . It remains disciplined by observable objects once G and the separability restriction are specified.

If an old vintage simply scales a firm's productive capacity, the correction above is enough. If instead obsolescence changes how firms learn, maintain knowledge, adopt technologies, or decide whether to exit, the effect cannot be removed from the observed size profile alone. The investment and exit policies would then vary with age, so an entrant that begins with mature knowledge need not remain distributed like mature firms. Evaluating codification would require specifying and disciplining those additional dynamics.

5.3 Elastic labor supply

The baseline holds aggregate employment fixed. What changes if a higher wage draws additional labor into production? Let $S \equiv M\bar{s}$ denote the economy's stock of organizational operating capacity. Firm life cycles continue to imply $S^c/S = y \cdot \ell$. The reason

is that every firm faces the same equilibrium wage. As shown in Lemma 1, a change in that wage scales current and future profits in the same way and therefore does not change firms' learning or exit policies. What changes is the amount of labor that can be combined with the additional organizational capacity. Conditional on the stock S of organization capacity and aggregate employment N , the production side of the economy can be written as

$$Y = S^{1-\nu} N^\nu.$$

Suppose labor supply is

$$N = \chi w^\varepsilon, \quad \varepsilon \geq 0, \quad (5.10)$$

where ε is the Frisch elasticity and $\chi > 0$ is a scale parameter.

Proposition 3. *With labor supply given by equation (5.10), the value of codification is*

$$V_\varepsilon = (y \cdot \ell)^{\omega(\varepsilon)}, \quad \omega(\varepsilon) \equiv \frac{(1-\nu)(1+\varepsilon)}{1+\varepsilon(1-\nu)}. \quad (5.11)$$

Appendix D.2 gives the derivation. The formula has a simple interpretation. When $\varepsilon = 0$, employment is fixed and $\omega(0) = 1 - \nu$, reproducing the baseline formula. As ε rises, codification raises the wage and draws additional labor into production. Firms can then combine the larger stock of organizational capacity with more labor rather than merely reallocating a fixed workforce. The exponent therefore rises. In the limit of perfectly elastic labor supply, $\omega(\varepsilon) \rightarrow 1$, the wage is fixed and output rises in proportion to organizational capacity. Thus, for $y\ell > 1$, allowing labor to respond increases the value of codification relative to the fixed-labor benchmark. The takeaway is that labor-supply elasticity changes the magnitude of the output response, but neither the two sufficient statistics nor their knowledge-based interpretation.

6 Conclusion

Codification allows knowledge accumulated through experience in one organization to become usable in another. Manuals, training systems, consulting, and franchising have long performed versions of this function. AI expands its potential because it can learn regularities directly from examples of practice, making capabilities portable even when the knowledge underlying them cannot be fully articulated. This paper develops a way to measure the economic value of that possibility.

We ask how much output would rise if entrants could begin with the organizational knowledge of mature firms. In a broad class of knowledge-based models of firm dynam-

ics, the answer depends on two observable statistics: mature firms' size relative to the average firm and their expected remaining lifespan relative to the expected lifespan of an entrant. Firm life cycles therefore provide a measure of the value of codification before direct evidence reveals how much organizational knowledge AI will actually make portable.

Applied to U.S. firm dynamics, our baseline estimate implies that complete codification approximately doubles aggregate output. Across industries, the value of codification varies widely and is only weakly related to current AI exposure. Exposure measures the work current systems can reach. Our sufficient statistic measures the economic stakes of codifying organizational knowledge. Assessing AI's potential requires both.

The estimates do not imply that AI will necessarily double output. They measure the value of eliminating the entrant--mature knowledge gap under the interpretation that firm life cycles reflect the accumulation of organizational knowledge. Their magnitude makes the remaining empirical questions more important. How much of the observed firm life cycle is generated by organizational knowledge rather than other age-dependent forces? How much of the knowledge that is not currently portable can AI codify, and how quickly can it make that knowledge portable across organizations? Answering these questions will require new evidence. The contribution of our sufficient statistic is to show, before that evidence exists, that the economic stakes are large.

Appendix

A Convergence

We want to show that, under the sequential implementation, the economy converges to the stationary equilibrium of Proposition 1: the entry distribution converges to h^m .

The key issue we need to tackle is that wages rise along the transition, so that the value function and policies change. In particular,

$$v_t(z, 1) = \max_{x \in [0,1]} \left\{ \pi(z, x, 1) + \beta \left(\frac{w_{t+1}}{w_t} \right)^{-\nu/(1-\nu)} \int v_{t+1}(z', 1) F(dz' \mid z, x) \right\},$$

so that wage growth $g_t = w_{t+1}/w_t$ enters as an adjustment to the discount factor. Along the transition, wages rise, so $g_t > 1$ for a while, and the operator at date t is some $\mathcal{T}_t \neq \mathcal{T}$. The question is whether this difference can prevent convergence.

Step 1: constant-growth economies. Take any fixed growth rate $g \geq 1$ and imagine an economy in which wages grow at that rate forever. Firms in that economy behave exactly as in a stationary economy with a lower discount factor, so everything we know about $\mathcal{T}$ applies to its operator $\mathcal{T}_g$: it has a unique invariant distribution, call it $h^m(g)$, and iterating from any starting point leads there. Then, higher anticipated wage growth makes the future less valuable, firms invest less, and so $\mathcal{T}_g$ moves a cohort's knowledge up by less than $\mathcal{T}$ does. Moreover, when g is close to 1, $h^m(g)$ is close to h^m .

Step 2: sandwich the sequence of entry distributions. Now take the true transition and suppose its wage growth rates all lie between 1 and some $\bar{g}$. At every date the true operator lies between $\mathcal{T}_{\bar{g}}$ and $\mathcal{T}$. If operators preserve the ordering of distributions — a cohort that starts with more knowledge ends with more, which holds when F is increasing in z and policies are monotone — then the true sequence of entry distributions is in between the two constant-growth ones: below it, the sequence generated by $\mathcal{T}_{\bar{g}}$; above it, the factual sequence $\mathcal{T}^a(h_0)$. By Step 1 the lower sequence converges to $h^m(\bar{g})$ and the upper one to h^m . So after enough time the true entry distribution lies between $h^m(\bar{g})$ and h^m .

Step 3: the sandwich closes as wages settle. Being in between $h^m(\bar{g})$ and h^m is not convergence to h^m , because the two differ. What closes the sandwich is that the bound can be tightened over time: if $g_t \rightarrow 1$, then for any late date we can redo Step 2 with a $\bar{g}$ close

to 1, the lower bound $h^m(\bar{g})$ moves up to h^m , and the true entry distribution is squeezed onto h^m . So the necessary condition wage growth per period goes to zero.

Step 4: condition for wages to settle. The wage is a function of the aggregate size index S_t , so wage growth goes to zero if S_t converges. S_t cannot exceed its stationary value in the counterfactual S^c because entrants never know more than mature firms. So if S_t rises over time, it converges, wage growth dies out, Step 3 applies, and the entry distribution converges to h^m — which forces S_t all the way to S^c . The one substantive condition is therefore that S_t rises along the transition. Two forces act on it: each cohort enters with more knowledge than the last, which pushes S_t up; and anticipated wage growth reduces investment, which pushes it down. The condition is that the first outweighs the second.

B AR(1) Illustration

This appendix gives a fully specified example satisfying Assumption 1. It has only two effective choices—continue or harvest—and a contracting AR(1) process for organizational capacity. The example is illustrative rather than calibrated. Its purpose is to make three distinct objects concrete: the convergence of a cohort to the mature distribution, the stationary comparison in Proposition 1, and the calendar-time adjustment after codification.

B.1 Knowledge dynamics and the mature distribution

Let $k \equiv z^{1/(1-\nu)}$ denote organizational operating capacity before the resources devoted to maintaining knowledge are deducted. If the firm chooses $x > 0$, it remains active next period with probability

$$\rho \min\{2x, 1\}, \quad \rho \in (0, 1).$$

Conditional on remaining active, its capacity follows

$$k' = (1 - \lambda)\mu + \lambda k + \epsilon', \quad \lambda \in (0, 1), \tag{B.1}$$

where ϵ' is an iid mean-zero shock with bounded support. Otherwise, the firm enters the absorbing exit state. As in the baseline model, choosing $x = 0$ gives the firm its undiminished current operating profit and certain exit before the next period.

Choices above $x = 1/2$ are dominated by $x = 1/2$: they produce less current output without increasing survival. Over $x \in [0, 1/2]$, the normalized objective is convex because

current profits are proportional to $k(1 - x)^{1/(1-\nu)}$ and continuation value is linear in x . The optimal choice is therefore either $x = 0$ or $x = 1/2$. Equivalently, ϕ equals one when the firm harvests and exits and one-half when it continues.

For concreteness, set

$$\nu = 0.5, \quad \beta = 0.99, \quad \rho = 0.9, \quad \lambda = 0.97, \quad \mu = 1,$$

and let ϵ' equal 0.003 or -0.003 with equal probability. Entrants begin with $k_0 = 0.05$. Under the policy $x = 1/2$, the normalized value of continuing forever is

$$\tilde{v}^c(k) = (1 - \nu)v^{\nu/(1-\nu)}2^{-1/(1-\nu)} \left[\frac{\mu}{1 - \beta\rho} + \frac{k - \mu}{1 - \beta\rho\lambda} \right] = 0.25(0.452 + 1.842k).$$

The invariant support of k is $[0.9, 1.1]$, and all states reached from k_0 lie in $(0, 1.1]$. Since $\tilde{v}^c(k) > 0.25k$ throughout this interval, continuing is optimal at every active state reached in the example. Firms nevertheless exit with probability $1 - \rho = 0.1$ each period.

The AR(1) makes Assumption 1 transparent. If two firms receive the same shock, the distance between their next-period capacities is λ times their current distance. Since $\lambda < 1$, the effect of the initial condition disappears and equation (B.1) has a unique invariant distribution h^m . This distribution is nondegenerate because firms continue to receive shocks, and its mean is μ . Because shutdown risk is independent of k , conditioning on survival does not alter the AR(1) transition. Thus h^m is also the unique invariant distribution of the survivor operator $\mathcal{T}$, and $\bar{\rho}^m = \rho$.

A factual cohort begins at k_0 , so its mean capacity and surviving mass at age a are

$$\mathbb{E}[k_a] = \mu + (k_0 - \mu)\lambda^a, \quad m_a = \rho^a.$$

Panel (a) of Figure 5 displays both sequences. Firms continue to move across knowledge states and some exit at every age, while the distribution conditional on survival converges to h^m .

B.2 Stationary comparison

Constant shutdown risk deliberately suppresses the lifespan margin: both entrants and mature firms have expected remaining lifespan $1/(1 - \rho) = 10$, so $\ell = 1$. The average

capacity in the factual stationary economy is

$$\bar{k} = (1 - \rho) \sum_{a=0}^{\infty} \rho^a \mathbb{E}[k_a] = \mu + (k_0 - \mu) \frac{1 - \rho}{1 - \rho\lambda} = 0.252. \quad (\text{B.2})$$

Because every active firm chooses $x = 1/2$, its size index is $s = 2^{-1/(1-\nu)}k$. The common factor cancels from the relative-size statistic, giving $y = \mu/\bar{k} = 3.97$. Complete codification replaces the entrant distribution with h^m , so every counterfactual cohort has the invariant distribution at every age. Proposition 1 then gives

$$V = (y\ell)^{1-\nu} = 3.97^{1/2} = 1.99.$$

The numerical values were chosen to make the stationary gain easy to read, not to reproduce the empirical estimate. Allowing shutdown risk to depend on capacity would add a lifespan margin but would obscure the closed-form mechanics.

B.3 Calendar-time convergence

The stationary gain does not reveal when it is realized. To see this, suppose that complete codification is introduced at the beginning of period zero. New entrants immediately draw from h^m , while factual incumbents retain their current knowledge and continue to follow equation (B.1). Let K_t^c denote the total capacity of active firms, omitting the common factor $2^{-1/(1-\nu)}$. Labor-market clearing determines the real wage in each period. Solving the Bellman equation backward from the counterfactual stationary value verifies that continuing remains optimal at every reachable state along both wage paths considered below. The firm distributions can therefore be propagated forward directly. Under complete codification,

$$K_t^c = \frac{\mu}{1 - \rho} - \frac{(\mu - k_0)(\rho\lambda)^{t+1}}{1 - \rho\lambda}.$$

The remaining capacity gap therefore shrinks by the factor $\rho\lambda$ each period. Persistence slows the learning of factual incumbents, while survival slows their replacement by mature entrants.

In the numerical example, $\rho\lambda = 0.873$. The capacity gap has a half-life of about 5.1 periods, and 90 percent of it disappears after about 17 periods. Output rises by 17 percent on impact, reaches 1.63 times factual output after five periods, and converges to almost twice factual output. The direct-codification path is shown in panel (b) of Figure 5.

These numbers come from the additional primitives of the example, not from the sufficient statistics. Indeed, changing λ and adjusting k_0 to keep equation (B.2) unchanged leaves y , ℓ , and the stationary value V unchanged, while changing the convergence rate $\rho\lambda$. Firm life-cycle moments identify the destination in Proposition 1; they do not identify the speed of the journey.

The same example also makes the sequential implementation in Section 3.3 concrete. If each cohort receives the conditional knowledge distribution attained after one period by the preceding cohort, the mean capacity of entrants in period t is $\mu + (k_0 - \mu)\lambda^t$. The entry distribution converges to h^m at rate λ , rather than jumping there at date zero. The resulting aggregate transition, also shown in panel (b), is slower than under direct codification. In this example the transition can be computed because the knowledge process and policy are fully specified and the policy remains fixed along the path. These are precisely the additional restrictions that the sufficient-statistic result avoids imposing.




Figure 5: A contracting AR(1) illustration. Panel (a) reports the factual cohort's mean operating capacity and surviving mass. Panel (b) reports output after direct codification and under the sequential one-period-ahead implementation.

C Linear Production and Differentiated Varieties

This appendix shows that the sufficient-statistic result does not rely on decreasing returns at the firm level. Suppose production is linear in labor but each firm supplies a differentiated variety. The productive role of additional firms then comes from the CES aggregator rather than from spreading labor across decreasing-returns technologies.

Let a competitive final-good producer combine the varieties supplied by active firms

according to

$$Y = \left(\int_{\mathcal{I}} q_i^{(\sigma-1)/\sigma} di \right)^{\sigma/(\sigma-1)}, \quad \sigma > 1,$$

where $\mathcal{I}$ is the set of active firms. A firm with knowledge z that chooses investment x produces its variety according to

$$q = z(1-x)n.$$

Thus, firm-level production is constant returns to labor. Normalize the CES price index to one and let w denote the real wage. Demand for variety i is

$$q_i = Y p_i^{-\sigma}.$$

C.1 The firm problem

Conditional on (z, x) , marginal cost is $w/[z(1-x)]$. A firm charges the constant markup $\mu \equiv \sigma/(\sigma-1)$, so

$$p_i = \frac{\mu w}{z_i(1-x_i)}.$$

Employment and operating profits are therefore

$$\begin{aligned} n_i &= Y \mu^{-\sigma} w^{-\sigma} [z_i(1-x_i)]^{\sigma-1}, \\ \pi_i &= \frac{w}{\sigma-1} n_i = \chi(w, Y) [z_i(1-x_i)]^{\sigma-1}. \end{aligned} \tag{C.1}$$

where $\chi(w, Y) \equiv Y \mu^{-\sigma} w^{1-\sigma} / (\sigma-1)$ is common to all firms.

The same homogeneity that makes policies independent of the real wage in the baseline applies here to the aggregate term $\chi(w, Y)$. In a stationary equilibrium, the Bellman equation is

$$v_\sigma(z; w, Y) = \max_{x \in [0,1]} \left\{ \chi(w, Y) [z(1-x)]^{\sigma-1} + \beta \int v_\sigma(z'; w, Y) F(dz' | z, x) \right\}.$$

Writing $v_\sigma(z; w, Y) = \chi(w, Y) \tilde{v}_\sigma(z)$ gives

$$\tilde{v}_\sigma(z) = \max_{x \in [0,1]} \left\{ [z(1-x)]^{\sigma-1} + \beta \int \tilde{v}_\sigma(z') F(dz' | z, x) \right\}.$$

The stationary investment and exit policy is consequently independent of w and Y . This is the differentiated-varieties analogue of Lemma 1: the stationary policy, and hence the

survivor operator and mature distribution, are unchanged by the general-equilibrium adjustment in w and Y .

Let $x_\sigma^*(z)$ denote the stationary policy in this economy and define the size index

$$s_\sigma(z) \equiv [z(1 - x_\sigma^*(z))]^{\sigma-1}.$$

Equation (C.1) shows that employment is proportional to $s_\sigma(z)$ at common aggregate conditions. Revenue is proportional to the same index. The empirical size statistic can therefore continue to be measured by employment.

C.2 Aggregation and the sufficient statistic

Define aggregate operating capacity as

$$S_\sigma \equiv \sum_{a=0}^{\infty} m_a \int s_\sigma(z) h_a(dz) = M\bar{s}_\sigma.$$

Using firms' optimal prices, the normalization of the CES price index requires

$$1 = \int_{\mathcal{I}} p_i^{1-\sigma} di = \mu^{1-\sigma} w^{1-\sigma} S_\sigma. \quad (\text{C.2})$$

Labor-market clearing and equation (C.1) imply

$$1 = Y\mu^{-\sigma} w^{-\sigma} S_\sigma. \quad (\text{C.3})$$

Combining equations (C.2) and (C.3) gives the equilibrium wage and output:

$$\begin{aligned} w &= \mu^{-1} S_\sigma^{1/(\sigma-1)}, \\ Y &= S_\sigma^{1/(\sigma-1)} = (M\bar{s}_\sigma)^{1/(\sigma-1)}. \end{aligned}$$

Under complete codification, entrants draw from the mature distribution. The stationary mass of firms is therefore $\bar{\ell}^m$ and their average size index is $\bar{s}_\sigma^m$. Hence

$$Y^c = (\bar{\ell}^m \bar{s}_\sigma^m)^{1/(\sigma-1)}.$$

Let $y \equiv \bar{s}_\sigma^m / \bar{s}_\sigma$ denote average employment among mature firms relative to average employment among all active firms, and let $\ell \equiv \bar{\ell}^m / \bar{\ell}$. Because the unit entry flow implies

$M = \bar{\ell}$, the value of complete codification is

$$V_\sigma \equiv \frac{Y^c}{Y} = (y\ell)^{1/(\sigma-1)}.$$

The two production structures therefore preserve the same sufficient-statistic logic: mature relative size and relative lifespan determine the stationary output gain. Only the aggregation elasticity changes. In the baseline, additional firms raise output by allowing labor to be spread across organizational units subject to decreasing returns. Here, additional firms raise output by expanding the set of varieties. Holding y and ℓ fixed, the gain weakens as σ rises and vanishes in the perfect-substitutes limit.

D Derivations for Section 6

This appendix derives the growth and elastic-labor formulas in Section 5.

D.1 Aggregate productivity growth

D.1.1 Productivity growth common to all firms

Under equation (5.6), optimal employment at date t is

$$n_t(z) = \left(\frac{\nu A_t}{w_t} \right)^{1/(1-\nu)} s(z).$$

Let $\tilde{w}_t = w_t/A_t$. In a balanced-growth equilibrium, $\tilde{w}_t = \tilde{w}$. Writing $v_t(z) = A_t \tilde{v}(z)$ in the firm's Bellman equation and dividing by A_t gives

$$\tilde{v}(z) = \max_{x \in [0,1]} \left\{ (1-\nu) \nu^{\nu/(1-\nu)} \tilde{w}^{-\nu/(1-\nu)} [z(1-x)]^{1/(1-\nu)} + \beta G \int \tilde{v}(z') F(dz' | z, x) \right\}.$$

This normalized problem is independent of calendar time. The condition $\beta G < 1$ keeps the discounted value finite. The resulting policy and survivor operator may depend on G but not on calendar time. The same normalized problem applies before and after codification because the counterfactual holds (A_t, G) fixed.

With fixed aggregate employment, labor-market clearing gives $1 = (\nu/\tilde{w}_t)^{1/(1-\nu)} S$. The firm first-order condition implies aggregate output $Y_t = w_t/\nu$, and therefore

$$\frac{Y_t}{A_t} = S^{1-\nu}.$$

Taking factual and counterfactual ratios proves equation (5.7).

D.1.2 Productivity growth embodied in new vintages

Consider the production function in equation (5.8) and define $\tilde{w} \equiv w_t/B_t$. With fixed labor and a stationary normalized age distribution, the labor-market-clearing equation below makes $\tilde{w}$ constant. Operating profits of an age- a firm are $(1-\nu)\nu^{\nu/(1-\nu)}B_t\tilde{w}^{-\nu/(1-\nu)}G^{-a/(1-\nu)}[z(1-x)]^{1/(1-\nu)}$. The guess $v_{a,t}(z) = B_t\tilde{w}^{-\nu/(1-\nu)}G^{-a/(1-\nu)}\tilde{v}_g(z)$ gives the normalized Bellman equation

$$\tilde{v}_g(z) = \max_{x \in [0,1]} \left\{ (1-\nu)\nu^{\nu/(1-\nu)} [z(1-x)]^{1/(1-\nu)} + \beta G^{-\nu/(1-\nu)} \int \tilde{v}_g(z') F(dz' | z, x) \right\}. \quad (\text{D.1})$$

The factor used to normalize value changes from one period to the next by $G^{-\nu/(1-\nu)}$ for every firm, regardless of its age. Vintage growth therefore changes the effective discounting of continuation values and can change the policy relative to the economy without growth. Equation (D.1) is nevertheless stationary and does not depend on age. Let $s_g(z) \equiv [z(1-x_g^*(z))]^{1/(1-\nu)}$ denote the resulting size index; below we suppress its g subscript. For a firm of age a , optimal employment is

$$n_{a,t}(z) = \left(\frac{\nu}{\tilde{w}} \right)^{1/(1-\nu)} G^{-a/(1-\nu)} s_g(z). \quad (\text{D.2})$$

Aggregating D.2 over the factual age distribution gives

$$1 = \left(\frac{\nu}{\tilde{w}} \right)^{1/(1-\nu)} \sum_{a=0}^{\infty} m_a G^{-a/(1-\nu)} \bar{s}_a = \left(\frac{\nu}{\tilde{w}} \right)^{1/(1-\nu)} S_g. \quad (\text{D.3})$$

Because the firm first-order condition is $\nu q_{a,t} = w_t n_{a,t}$, aggregate output is $Y_t = w_t/\nu$. Equation (D.3) then implies

$$\frac{Y_t}{B_t} = S_g^{1-\nu}. \quad (\text{D.4})$$

Under complete codification, every cohort begins with h^m . Under the separability restriction, its conditional knowledge distribution remains h^m and its mass at age a is $(\bar{\rho}^m)^a$. Counterfactual labor-market clearing therefore uses

$$S_g^c = \sum_{a=0}^{\infty} (\bar{\rho}^m)^a G^{-a/(1-\nu)} \bar{s}^m = \frac{\bar{s}^m}{1 - \bar{\rho}^m G^{-1/(1-\nu)}}.$$

Taking factual and counterfactual ratios in equation (D.4) therefore gives equation (5.9).

D.2 Elastic labor supply

Let $S = M\bar{s}$ denote the aggregate stock of organizational operating capacity. Aggregating across firms gives the production frontier

$$Y = S^{1-\nu} N^{\nu}.$$

The corresponding wage is

$$w = \nu S^{1-\nu} N^{\nu-1}.$$

Combining this expression with $N = \chi w^{\varepsilon}$ yields

$$N^{1+\varepsilon(1-\nu)} = \chi \nu^{\varepsilon} S^{\varepsilon(1-\nu)},$$

and therefore

$$N = (\chi \nu^{\varepsilon})^{\frac{1}{1+\varepsilon(1-\nu)}} S^{\frac{\varepsilon(1-\nu)}{1+\varepsilon(1-\nu)}}.$$

Substituting employment into the production frontier gives

$$Y = C(\chi, \nu, \varepsilon) S^{\omega(\varepsilon)}, \quad \omega(\varepsilon) = \frac{(1-\nu)(1+\varepsilon)}{1+\varepsilon(1-\nu)},$$

where the constant $C(\chi, \nu, \varepsilon)$ does not depend on organizational capacity and is therefore common to the factual and counterfactual economies. Because complete codification implies $S^c/S = y\ell$, taking the output ratio gives equation (5.11).

E Additional empirical results

E.1 Aggregate firm and establishment estimates

Table 1 reports the inputs to equation (3.1) when the organizational unit is a firm or an establishment. The lifespan ratios are similar across the two units. The difference in the value of codification comes primarily from mature firms' expansion across establishments.

Unit	Lifespan			Employment			
	Average	Mature remaining	ℓ	Average	Mature	y	V
Firm	10.4	23.4	2.24	24.6	174.3	7.09	2.00
Establishment	10.4	22.4	2.16	18.3	49.4	2.69	1.55

Table 1: Aggregate estimates by empirical unit

E.2 Industry estimates by empirical unit

Table 2 reports the screened industry sample for firms and establishments side by side. As in Section 4, the sample requires at least 50 mature-firm deaths in 2023. The firm columns correspond to our baseline unit of analysis. In the establishment columns, the maturity clock resets when a location opens; these estimates therefore capture knowledge accumulated within a location but not capabilities retained by a parent firm and applied across locations. Reporting both units in the same table makes this distinction transparent and avoids repeating the industry and exposure columns.

Table 2: Industry estimates by empirical unit

Code	Industry	Firm			Establishment			Exp.
		V	y	ℓ	V	y	ℓ	
236	Construction of Bldgs	2.03	7.10	2.38	1.77	4.07	2.39	0.31
237	Heavy/Civil Eng Const	1.71	4.07	2.13	1.47	2.13	2.20	0.22
238	Specialty Trade Const	1.71	4.62	1.85	1.57	2.92	2.08	0.22
311	Food Mfg	1.90	5.93	2.19	1.59	2.82	2.28	0.19
321	Wood Prod Mfg	1.64	4.27	1.68	1.37	2.10	1.69	0.19
323	Printing	1.45	3.07	1.45	1.27	1.92	1.36	0.35
327	Nonmetallic Mineral	1.50	2.91	1.76	1.23	2.01	1.15	0.25
332	Fabricated Metal	1.47	2.70	1.74	1.32	1.77	1.73	0.27
333	Machinery Mfg	1.66	2.87	2.66	1.43	1.89	2.19	0.31
339	Misc Mfg	1.68	4.61	1.74	1.43	2.45	1.70	0.31
423	Wholesale: Durable	1.72	4.33	2.02	1.39	1.67	2.20	0.41
424	Wholesale: Nondurable	1.88	5.22	2.38	1.45	2.10	2.09	0.38
425	Wholesale: Electronic	2.02	8.29	2.03	1.39	1.97	1.89	0.49
441	Motor Vehicle Dealers	1.54	3.62	1.57	1.23	1.74	1.32	0.30
442	Furniture & Electronics	1.67	3.50	2.23	1.10	1.04	1.40	0.39
444	Bldg Material Dealers	1.43	2.60	1.61	0.97	0.67	1.30	0.32
445	Food & Bev Stores	2.25	15.09	1.70	1.34	1.75	1.86	0.25
446	Health/Personal Care	2.02	9.09	1.83	1.09	1.10	1.29	0.38
447	Gasoline Stations	1.97	10.40	1.45	1.07	1.07	1.22	0.26
448	Clothing Stores	2.13	11.27	1.82	1.07	1.30	1.02	0.36
451	Other Retailers	1.87	5.90	2.21	1.23	1.17	1.92	0.35

Continued on next page

Table 2 – continued

Code	Industry	Firm			Establishment			Exp.
		V	y	ℓ	V	y	ℓ	
484	Truck Transportation	2.57	12.99	3.38	1.84	3.51	3.26	0.31
511	Publishing	1.66	3.04	2.47	1.27	1.30	2.02	0.63
522	Credit Intermediation	2.04	3.12	5.56	1.32	1.16	2.63	0.57
524	Insurance Carriers	1.94	9.62	1.46	1.49	2.99	1.63	0.58
531	Real Estate	1.71	3.92	2.20	1.34	1.72	1.87	0.39
532	Rental & Leasing	2.11	8.83	2.24	1.21	2.05	1.05	0.35
541	Prof/Sci/Tech Svcs	2.11	10.32	1.92	1.52	2.85	1.86	0.55
551	Mgmt of Companies	1.55	1.79	3.21	1.61	2.17	3.08	0.53
561	Admin & Support	2.09	8.58	2.24	1.29	1.38	2.02	0.28
611	Educational Svcs	2.20	5.64	4.16	2.24	5.78	4.39	0.38
621	Ambulatory Health	1.69	7.54	1.09	1.17	1.89	1.01	0.37
623	Nursing/Resid Care	1.80	2.94	3.54	1.63	2.80	2.54	0.22
624	Social Assistance	1.87	3.46	3.51	1.54	1.95	2.90	0.24
713	Amusement/Gambling	1.81	2.97	3.62	1.54	1.97	2.84	0.24
721	Accommodation	1.78	6.05	1.65	1.19	1.55	1.29	0.23
722	Food Svcs & Drinking	1.99	9.04	1.74	1.18	1.25	1.53	0.17
811	Repair & Maint	1.53	3.34	1.62	1.18	1.25	1.54	0.24
812	Personal/Laundry	1.71	4.69	1.83	1.33	1.79	1.72	0.23
813	Religious/Civic Org	1.43	1.77	2.34	1.39	1.63	2.33	0.42

Notes: The table reports the value of codification V , the mature–average size ratio y , and the mature–entrant lifespan ratio ℓ for each empirical unit. LLM exposure is the occupational measure of Eloundou et al. (2024). The table includes industries with at least 50 mature-firm deaths in 2023, as in the main industry analysis. A dash indicates that the estimate is unavailable for that unit.

References

- Daron Acemoglu. The simple macroeconomics of AI. *Economic Policy*, 40(121):13–58, 2025.
- Daron Acemoglu and Pascual Restrepo. The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review*, 108(6):1488–1542, 2018. doi: 10.1257/aer.20160696.
- Philippe Aghion, Benjamin F. Jones, and Charles I. Jones. Artificial intelligence and economic growth. Working Paper 23928, National Bureau of Economic Research, 2017.
- David Argente, Sara Moreira, Ezra Oberfield, and Venky Venkateswaran. *Scalable expertise*. National Bureau of Economic Research, 2025.
- Linda Argote. *Organizational Learning: Creating, Retaining and Transferring Knowledge*. Springer, New York, 2 edition, 2013.
- Andrew Atkeson and Patrick J. Kehoe. Modeling and measuring organization capital. *Journal of Political Economy*, 113(5):1026–1053, 2005. doi: 10.1086/431289.
- David Autor. Applying ai to rebuild middle class jobs. Technical report, National Bureau of Economic Research, 2024.
- Tania Babina, Alex Xi He, and Renhao Jiang. Canaries in the gold mine: Early productivity gains from artificial intelligence creating organization capital. Working paper, 2026.
- Byong-Hyong Bahk and Michael Gort. Decomposing learning by doing in new plants. *Journal of Political Economy*, 101(4):561–583, 1993.
- Isaac Baley and Laura Veldkamp. Bayesian learning. In Rüdiger Bachmann, Giorgio Topa, and Wilbert van der Klaauw, editors, *Handbook of Economic Expectations*, pages 717–748. Academic Press, Amsterdam, 2023. doi: 10.1016/B978-0-12-822927-9.00032-X. URL <https://doi.org/10.1016/B978-0-12-822927-9.00032-X>.
- Isaac Baley and Laura L. Veldkamp. *The Data Economy: Tools and Applications*. Princeton University Press, Princeton, NJ, 2025. ISBN 9780691256726. doi: 10.1515/9780691256740. URL <https://doi.org/10.1515/9780691256740>.
- C. Lanier Benkard. Learning and forgetting: The dynamics of aircraft production. *American Economic Review*, 90(4):1034–1054, 2000.

- Martin Beraja and Nathan Zorzi. Inefficient automation. *Review of Economic Studies*, 92(1): 69–96, 2025. doi: 10.1093/restud/rdae019.
- Nicholas Bloom, Raffaella Sadun, and John Van Reenen. Management as a technology? Working Paper 22327, National Bureau of Economic Research, 2016. Revised October 2017.
- 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 R. Raymond. Generative AI at work. *Quarterly Journal of Economics*, 140(2):889–942, 2025. doi: 10.1093/qje/qjae044.
- Francisco J. Buera and Ezra Oberfield. The global diffusion of ideas. *Econometrica*, 88(1): 83–114, 2020. doi: 10.3982/ECTA14044.
- Aaron Chatterji, Daniel Rock, and Eduard Talamas. The human-machine knowledge spiral. *arXiv preprint arXiv:2606.29227*, 2026.
- Robin Cowan and Dominique Foray. The economics of codification and the diffusion of knowledge. *Industrial and Corporate Change*, 6(3):595–622, 1997. doi: 10.1093/icc/6.3.595.
- Eric D. Darr, Linda Argote, and Dennis Epple. The acquisition, transfer, and depreciation of knowledge in service organizations: Productivity in franchises. *Management Science*, 41(11):1750–1762, 1995. doi: 10.1287/mnsc.41.11.1750.
- Joel M. David, Hugo A. Hopenhayn, and Venky Venkateswaran. Information, misallocation, and aggregate productivity. *Quarterly Journal of Economics*, 131(2):943–1005, 2016. doi: 10.1093/qje/qjw006.
- Tom Davidson, Basil Halperin, Thomas Houlden, and Anton Korinek. When does automating ai research produce explosive growth? feedback loops in innovation networks. Technical report, National Bureau of Economic Research, 2026.
- Avinash K. Dixit and Joseph E. Stiglitz. Monopolistic competition and optimum product diversity. *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. doi: 10.1126/science.adj0998.
- Marcela Eslava, John Haltiwanger, and Nicolás Urdaneta. The size and life-cycle growth of plants: The role of productivity, demand, and wedges. *Review of Economic Studies*, 91(1):259–300, 2024. doi: 10.1093/restud/rdad029.
- Lucia Foster, John Haltiwanger, and Chad Syverson. Reallocation, firm turnover, and efficiency: Selection on productivity or profitability? *American Economic Review*, 98(1):394–425, 2008.
- Robert M. Grant. Toward a knowledge-based theory of the firm. *Strategic Management Journal*, 17(S2):109–122, 1996.
- Hugo A. Hopenhayn. Firms, misallocation, and aggregate productivity: A review. *Annual Review of Economics*, 6:735–770, 2014.
- Chang-Tai Hsieh and Peter J. Klenow. Misallocation and manufacturing TFP in china and india. *Quarterly Journal of Economics*, 124(4):1403–1448, 2009.
- Chang-Tai Hsieh and Peter J. Klenow. The life cycle of plants in india and mexico. *Quarterly Journal of Economics*, 129(3):1035–1084, 2014. doi: 10.1093/qje/qju014.
- Enrique Ide and Eduard Talamàs. Artificial intelligence in the knowledge economy. *Journal of Political Economy*, 133(12):3762–3800, 2025. doi: 10.1086/737233.
- Enrique Ide and Eduard Talamàs. The turing valley: How ai capabilities shape labor income. *arXiv preprint arXiv:2408.16443*, 2026.
- Benjamin F. Jones. Artificial intelligence in research and development. Working Paper 34312, National Bureau of Economic Research, 2025.
- Charles I. Jones and Christopher Tonetti. Past automation and future AI: How weak links tame the growth explosion. Working Paper, Stanford University and NBER, 2026.
- Boyan Jovanovic and Yaw Nyarko. A bayesian learning model fitted to a variety of empirical learning curves. *Brookings Papers on Economic Activity. Microeconomics*, 1995:247–305, 1995.
- Boyan Jovanovic and Rafael Rob. The growth and diffusion of knowledge. *The Review of Economic Studies*, 56(4):569–582, 1989.

- Réka Juhász, Sho Sakabe, and David E. Weinstein. Codification, technology absorption, and the globalization of the industrial revolution. Working paper, Columbia University, 2025. November 2025.
- Steven D. Levitt, John A. List, and Chad Syverson. Toward an understanding of learning by doing: Evidence from an automobile assembly plant. *Journal of Political Economy*, 121(4):643–681, 2013.
- Jr. Lucas, Robert E. On the size distribution of business firms. *Bell Journal of Economics*, 9(2):508–523, 1978.
- Jr. Lucas, Robert E. and Benjamin Moll. Knowledge growth and the allocation of time. *Journal of Political Economy*, 122(1):1–51, 2014. doi: 10.1086/674363.
- Erzo G. J. Luttmer. Selection, growth, and the size distribution of firms. *Quarterly Journal of Economics*, 122(3):1103–1144, 2007. doi: 10.1162/qjec.122.3.1103.
- Joel Mokyr. *The Gifts of Athena: Historical Origins of the Knowledge Economy*. Princeton University Press, Princeton, 2002.
- Ikujiro Nonaka. A dynamic theory of organizational knowledge creation. *Organization Science*, 5(1):14–37, 1994. doi: 10.1287/orsc.5.1.14.
- Jesse Perla and Christopher Tonetti. Equilibrium imitation and growth. *Journal of Political Economy*, 122(1):52–76, 2014. doi: 10.1086/674362.
- Billy Perrigo. Why Google bid \$10 million for a failed airline’s data. *Time*, August 25, 2026, 2026. URL <https://time.com/article/2026/08/25/google-spirit-airlines-ai-data-RL/>.
- Michael Polanyi. *The Tacit Dimension*. Doubleday, Garden City, NY, 1966.
- Edward C. Prescott and Michael Visscher. Organization capital. *Journal of Political Economy*, 88(3):446–461, 1980.
- David Silver and Richard S. Sutton. Welcome to the era of experience. Preprint of a chapter forthcoming in *Designing an Intelligence*, MIT Press, 2025. URL <https://storage.googleapis.com/deepmind-media/Era-of-Experience%20/The%20Era%20of%20Experience%20Paper.pdf>.

- Gabriel Szulanski. Exploring internal stickiness: Impediments to the transfer of best practice within the firm. *Strategic Management Journal*, 17(S2):27–43, 1996. doi: 10.1002/smj.4250171105.
- Frederick Winslow Taylor. *The Principles of Scientific Management*. Harper & Brothers, New York, 1911.
- David J. Teece. The market for know-how and the efficient international transfer of technology. *Annals of the American Academy of Political and Social Science*, 458:81–96, 1981.
- U.S. Census Bureau. Business dynamics statistics: Firm and establishment age, 2023. Business Dynamics Statistics data files, 2025. Accessed July 2026.
- Sidney G. Winter, Gabriel Szulanski, Dimo Ringov, and Robert J. Jensen. Reproducing knowledge: Inaccurate replication and failure in franchise organizations. *Organization Science*, 23(3):672–685, 2012. doi: 10.1287/orsc.1110.0663.
- Udo Zander and Bruce Kogut. Knowledge and the speed of the transfer and imitation of organizational capabilities: An empirical test. *Organization Science*, 6(1):76–92, 1995.