

Why Don't Old Firms Do New Things?*

Nicolas Crouzet[†]

Zhiguo He[‡]

Victor Lyonnet[§]

Yueran Ma[¶]

March 2026

Abstract

Since at least [Schumpeter \(1942\)](#), new firms have been widely viewed as the driving force for implementing new technologies. Yet limited economics research explains when and why new technologies require new firms. We examine the view that old firms struggle especially when new technologies lead to different organizational workstyles (based on occupation composition and their corresponding workstyles), which may clash with their existing business model. In the data, young firms grow significantly faster than old firms when new technologies in an industry generate greater workstyle changes (due to the types of workers they require), whereas the advantage of young firms is not related to the sheer volume of new technologies (e.g., the number of all or important patents). We also find that new technologies associated with greater workstyle changes are more likely to be implemented by young firms, as measured using AI assessment of Wikipedia titles. These results highlight the role of organizational frictions in shaping companies' adaptability, and provide new perspectives for the [Coase \(1937\)](#) boundary of the firm question. Investment in entrepreneurship is especially important when innovations alter organizational workstyles.

*We are grateful to Geneve Bullo, Jingyi Jia, Howard Guankun Li, Haomin Qin, Raymon Yue, and Mengdi Zhang for extraordinary research assistance. We thank Nick Bloom, Ed Glaeser, Oliver Hart, Greg Kaplan, Steve Kaplan, Andrei Shleifer, Chad Syverson, Rob Vishny, and seminar participants at Boston University and e61 for insightful comments. First draft: November 2025.

[†]Northwestern University Kellogg School of Management (n-crouzet@kellogg.northwestern.edu).

[‡]Stanford University Graduate School of Business and NBER (hezhg@stanford.edu).

[§]University of Michigan Ross School of Business (vlyonnet@umich.edu).

[¶]University of Chicago Booth School of Business and NBER (yueran.ma@chicagobooth.edu).

I Introduction

Nearly one century ago, [Schumpeter \(1942\)](#) highlighted the importance of new technologies for economic progress. In his view, the implementation of new technologies relies crucially on entrepreneurship, and the process of creative destruction involves *new firms* overtaking existing ones. This view has been widely held among academics, policymakers, and the general public ever since ([Decker et al., 2014](#); [Draghi, 2024](#)), and entrepreneurship is a core topic of economics research today. However, some studies have found that existing firms contribute substantially to technological advancement ([Garcia-Macia, Hsieh, and Klenow, 2019](#); [Braguinsky et al., 2024](#)), and we can name plenty of new technologies that are successfully implemented by incumbents, such as polymers, smartphones, and cloud computing (which we discuss in more detail in Section II). The necessity of new firms for new technologies is not entirely self-evident.

When—and why—do new technologies require new firms to implement? Although this question is fundamental for understanding the essence of creative destruction, economics research offers limited theoretical guidance and systematic analyses. Influential models of creative destruction such as [Aghion and Howitt \(1992\)](#) do not address whether an entrant using a new technology is a new firm or an existing firm. Meanwhile, research motivated by the classic boundary of the firm question following [Coase \(1937\)](#) has focused on vertical integration ([Williamson, 1971](#); [Grossman and Hart, 1986](#); [Hart and Moore, 1990](#)), or “which transactions can be implemented within a firm?” Yet another aspect of the Coasian problem has received much less attention: “which ideas can be implemented within a firm?”

Case studies in management and reflections by managers often come to the view that old firms can implement new technologies if they are compatible with existing organizational processes and priorities, but struggle if they require new ones ([Christensen, 1997](#); [Gerstner, 2002](#); [O’Reilly and Tushman, 2021](#)). As Microsoft CEO Satya Nadella remarked at Chicago Booth’s 125th anniversary celebration: “this is one of the foundational challenges—when the new thing comes not only as a technology challenge but also a business model challenge, most companies can’t make it.” This view also echoes observations about the limits of organization in [Arrow \(1964\)](#), who postulates that organizations form “codes... in accordance with the best expectations of the firm’s creation,” which will be “modified only slowly over time,” in which case new technologies that require a different set of codes can be difficult for incumbents compared to new firms that start from scratch. The incompatibility between new technologies and existing organizational codes is then central to the need for new organizations.

In this paper, we formalize this hypothesis theoretically, and conduct extensive empirical analyses to document its relevance for the performance of young vs old firms in the face of new technologies. Our primary challenge—conceptually and empirically—is finding a way to capture the extent to which new

technologies require new organizational codes. Our entry point is to study how different technologies require different occupations to implement, which are associated with different workstyles. For example, in the car industry, technological advancements in software have proven especially challenging for old incumbents—they require software engineers who emphasize flexibility and creativity, which differ from the traditional core values of hardware engineers (whereas advancements in hardware were easy to integrate into their existing systems). Numerous news reports attest to the difficulty for legacy car makers to adapt to the different workstyles associated with software (George, 2022; Davis, 2023). Accordingly, we present a model that lays out how new technologies associated with larger changes in workstyles can lead to more severe organizational frictions, especially in old firms. We then develop a new and theoretically-grounded measure of workstyle changes induced by new technologies, which allows us to perform large-scale empirical tests about the performance of young vs old firms.

In our model, firms initially operate an existing business model A , which has a continuum of business units, each having a manager who is in charge of hiring its workers. The occupation composition among workers in A is pinned down by the productivity of the occupations given the technology associated with that business model. Later, a new business model B emerges, which has its own occupation-specific productivity that determines the desired occupation composition for B . Each occupation has a set of workstyles (i.e., characteristics it emphasizes), so the overall (employment-weighted average) workstyle under the new business model B can differ from that under the existing business model A .

When differences in workstyle are larger, there is a higher probability that workers in business model B need to solve a problem differently from what business model A is used to. This leads to a “conflict,” in which case managers in business model B need to justify their approach, in light of existing rules laid down by business model A that codify how to solve problems according to business model A ’s past experiences. Doing so takes time and reduces the productivity of business model B , which shrinks its optimal size. Older firms have more existing rules, so the time and productivity losses are greater. Therefore, they end up with a smaller operation of business model B . Accordingly, they grow more slowly when new technologies require larger changes in workstyle.

To validate that older firms indeed have more rules and more meetings that waste time, we process employee reviews from Revelio to measure the extent to which employees mention the presence of rules and the excess of meetings. Across different companies at a given point in time, we find that the fraction of reviews mentioning the presence of rules and the fraction of reviews mentioning the excess of meetings are positively correlated with firm age.

We build on the model to develop a new measure of workstyle changes in an industry induced by new technologies, which can be applied systematically across industries and over time. We proceed in two steps.

First, we predict future occupation composition in light of new technologies in an industry following [Kogan et al. \(2024\)](#). The key idea is to capture the exposure of the occupations' routine and nonroutine tasks to new technologies measured through patents, based on the textual similarity between occupation task descriptions and patent text. As [Kogan et al. \(2024\)](#) show, when new technologies in an industry are more similar to an occupation's routine (nonroutine) tasks, substitutability (complementarity) is stronger, and the employment of the occupation falls (rises) going forward (using data on employment by industry and occupation from the Bureau of Labor Statistics). Correspondingly, the similarity between patents in an industry-year and an occupation's tasks allows us to predict the occupation's future employment in the industry. Second, once we have future industry-level occupation composition predicted by new technologies, we can obtain the associated (employment-weighted average) overall future workstyle in the industry, relative to the current overall workstyle (using data on the workstyle associated with each occupation from O*NET). The predicted industry-level workstyle change (over the next five years) induced by the new technologies is the key independent variable in our empirical analyses. The reliance of the measurement on BLS and O*NET data restricts our sample period to 2003 onward.

We use three sets of data to test the core hypothesis that young firms grow more than old firms when new technologies are associated with greater changes in workstyle. First, we investigate venture capital (VC) investment volume in an industry-year, which captures the forward-looking valuation of startups that can reflect their growth potential. We find that a one standard deviation increase in the predicted industry-level workstyle change is associated with around 0.4 log points higher VC investment in the industry, controlling for the log of total market capitalization of Compustat firms in the industry to capture other factors that can affect the prospects of firms in the industry. Importantly, the volume of new technologies per se—measured as the total number of patents, breakthrough patents ([Kelly et al., 2021](#)), creative patents ([Kalyani, 2025](#)), or rapidly evolving patents ([Bowen, Frésard, and Hoberg, 2023](#)) in the industry—does not have a significant relationship with VC investment, or affect the coefficient on workstyle change—our key variable of interest.

Second, we investigate variation among Compustat firms, including their valuation, as well as realized sales growth and employment growth over the next five years. Specifically, we regress firm-level Tobin's Q or subsequent growth on log firm age interacted with the predicted industry-level workstyle change (based on the new technologies in the industry as before). We find that young firms have significantly higher valuation and realized future growth relative to old firms when the technology-induced workstyle change is larger. Meanwhile, the number of total patents, breakthrough patents, creative patents, or rapidly evolving patents in the industry is not associated with significantly different outcomes for young vs old firms. In other words, old firms are not invariably incapable when the amount of new technologies is

greater, or when the new technologies are more important, influential, or creative. We also show that the effects of firm age are not due to age being correlated with size, or due to older firms being profitable so they suffer more existing businesses being replaced and cannibalized (Arrow, 1962).

Third, we investigate the population of firms in the Census Business Dynamics Statistics (BDS) dataset, which has broad coverage (although many young firms in this case may not be aspirational entrepreneurship). BDS provides total employment of firms by age group: 0, 1-5, 6-10, 11-15, 16-20, 21-25, and the remaining age groups cannot be consistently defined over our sample period because firms' precise age is unknown if they are born before 1976. Therefore we restrict to firms with age groups between 1-5 and 16-20 (which become age groups 6-10 to 21-25 five years later). We find that the young age groups grow significantly faster than the older ones when the technology-induced workstyle change is larger. Again, the number of patents is not associated with significantly different outcomes for young vs old firms.

Furthermore, we show that the main results hold if we measure technology-induced workstyle change using Wikipedia titles instead of patents. Specifically, we capture new technologies using Wikipedia titles on technological inventions, and predict future occupation employment in an industry based on the exposure of the routine and nonroutine tasks, measured using the similarity between occupation task descriptions and Wikipedia titles' text (rather than patent text). All of our main results hold using predicted workstyle change based on Wikipedia titles instead. Additionally, we can estimate the predicted workstyle change associated with each Wikipedia title, and independently query LLM about whether the technological invention in a given Wikipedia title is most successfully implemented by an old or young (less than ten years old) firm.¹ We find that technological inventions associated with larger workstyle change (due to their predicted effects on occupation composition) are significantly more likely to be implemented by young firms according to the independent LLM query.

Taken together, these results support the importance of organizational frictions in shaping companies' adaptability. They show that new technologies per se do not necessarily challenge old firms. Incumbents are not ubiquitously incompetent in light of new technologies due to universal lack of learning or fear of cannibalization. However, when new technologies require changes in organizational workstyles, old firms struggle and young firms rise, in which case venture investment and government policies that facilitate entrepreneurship can be especially useful. Both our model and our empirical analyses on the growth of young vs old firms focus on the extent to which companies can *implement* new technologies, and we take the emergence of new technologies as given (e.g., artificial intelligence emerges from decades of research spanning academic and commercial domains, and information technology has roots in both publicly and

¹This exercise is much harder for patents which are more micro and technical, so we cannot easily ask LLM to determine whether the technologies associated with an individual patent are implemented by an old or young firm.

privately funded research). Once the new technology emerges, in principle any company can work to use it. The implementation of new technologies is foundational for the process of creative destruction and for business dynamism.

Literature Review Our work relates to three sets of literature. First, we contribute to research on creative destruction and the implementation of new technologies. We perform systematic analyses of the performance of young vs old firms in the face of new technologies, and highlight the role of organizational challenges. Prior work often sidesteps whether an entrant using a new technology is a new firm or an incumbent firm (Aghion and Howitt, 1992). Some analyses maintain that young firms are important for technological advancement (Decker et al., 2014; Loderer, Stulz, and Waelchli, 2017; Acemoglu et al., 2018; Draghi, 2024; Ewens and Marx, 2024), while others show that incumbents contribute meaningfully (Garcia-Macia, Hsieh, and Klenow, 2019; Cohen, Gurun, and Nguyen, 2022; Braguinsky et al., 2024). A recent study by Bowen, Frésard, and Hoberg (2023) shows that a greater quantity of “rapidly evolving” patents is associated with startups exiting through initial public offering—thus remaining independent—instead of selling out. Caskurlu, Hoberg, and Phillips (2024) focus on firm size and show that a rise of new technologies highly correlated across multiple industries is associated with faster growth by small firms. Our focus is on the importance of organizational frictions for understanding the implications of technological change for young vs old firms.

Second, we contribute to research on the nature of the firm. Many studies since Coase (1937) have focused on vertical integration (Williamson, 1971; Grossman and Hart, 1986; Hart and Moore, 1990; Aghion and Tirole, 1994; Anton and Yao, 1995), using insights of incomplete contracts and property rights,² but other considerations can also shape the activities that firms engage in (Atalay, Hortaçsu, and Syverson, 2014). Our work concerns another dimension of the boundary of the firm: which ideas can be implemented within a firm? We highlight challenges incumbents face that are not directly related to vertical integration, and point to a different set of mechanisms that restrict what a firm can do. In this regard, the most closely related work is Hart and Holmstrom (2010) on firm scope, which builds on the framework of Hart and Moore (2008) where authority—rather than bargaining à la Hart and Moore (1990)—plays a central role in resolving ex post conflicts. The key concept in Hart and Holmstrom (2010) is shading, an action motivated by one party’s grievance that imposes negative externalities on others, which becomes more severe with the integration of different business lines. This mechanism resembles the frictions between existing and new business models in our framework. However, instead of shading,

²Strictly speaking, the core takeaway for Grossman and Hart (1986), which highlights that asset ownership can be viewed as residual control rights, can apply to lateral integration too. Rajan and Zingales (2001) is an example which studies firm’s organization structure based on the framework of property rights, showing that flat hierarchies feature more distinctive technologies (than steep ones).

we focus on conflicts among workers and managers when the existing business model insists on legacy rules that emerge endogenously from the organization's dynamic structure. In the appendix, we show that our key prediction continues to hold under a more traditional setting in which conflicts arise between different business divisions under the same headquarters. This formulation draws on the insight of [Hart and Moore \(2008\)](#), where ex post deviations of realized payoffs from the preferred outcome specified in an ex ante contract lead to shading by division managers.³

Third, we contribute to research on organizational economics. The extent to which incumbent firms can implement new technologies is a classic question in management, which has produced a large volume of qualitative assessments and cases ([Penrose, 1959](#); [Nelson and Winter, 1985](#); [Tushman and Anderson, 1986](#); [Henderson and Clark, 1990](#); [March, 1991](#); [Leonard-Barton, 1992](#); [Teece, Pisano, and Shuen, 1997](#); [Christensen, 1997](#); [O'Reilly and Tushman, 2021](#)). We draw on their insights to develop formal theoretical modeling and systematic empirical analyses. This research question is also related to [Holmstrom \(1989\)](#), who investigates why innovation is typically undertaken by smaller firms using the complete contracting framework of [Holmstrom and Milgrom \(1991\)](#), unlike the incomplete contracting approach in the firm-boundary literature discussed above. The key idea is that combining hard-to-measure activities (e.g., innovation) with easy-to-measure activities (e.g., routine tasks) is costly due to the resulting misallocation of attention and effort. Another set of models study coordination and adaptation in organizations, emphasizing how firms use local information in different settings ([Milgrom and Roberts, 1992](#));⁴ our focus is organizational conflicts even without informational frictions. Finally, our modeling relates to work on knowledge hierarchy ([Garicano, 2000](#); [Garicano and Rossi-Hansberg, 2004](#)). Our mechanism incorporates the importance of managerial time, but our focus is the time loss from the conflict between business models, rather than time saving from the knowledge hierarchy.

The rest of the paper proceeds as follows. Section II presents stylized facts about the extent to which new technologies are implemented by new vs existing firms. Section III provides a model to motivate and organize our main empirical analyses. Section IV explains the data and measurement. Sections V and VI show the main empirical results. Section VII concludes.

³Our paper is also related to the literature on internal capital markets following the incomplete contracting tradition ([Shleifer and Vishny, 1989](#); [Meyer, Milgrom, and Roberts, 1992](#); [Stein, 1997](#); [Scharfstein and Stein, 2000](#); [Rajan, Servaes, and Zingales, 2000](#)), which typically focuses on the conflict between division managers and firm headquarters. In comparison, the main focus of our model is conflicts arising from workers in different business units.

⁴A long strand of the organizational economics literature emphasizes the role of local information and corresponding agency frictions across layers of decision makers, including formal vs real authority ([Aghion and Tirole, 1995, 1997](#)), corporate culture ([Gorton and Zentefis, 2024](#)), the design of functional units ([Qian, Roland, and Xu, 2006](#); [Dessein, Garicano, and Gertner, 2010](#)), and adaptation to changing environments ([Dessein and Santos, 2006](#)).

II New Technologies and New Firms: Some Stylized Facts

In this section, we present stylized facts to show that significant technological inventions are not invariably implemented by young firms. Old firms can implement new technologies, but not always. In other words, some new technologies may favor young firms, but others can be done by old firms. The fact that the implementation of new technologies is split between young and old firms motivates our subsequent analyses that investigate the circumstances where new technologies favor new firms.

II.A Technological Inventions Implemented by Young vs Old Firms

To obtain an intuitive sense of the extent to which new technologies are implemented by new firms, we proceed as follows. First, we collect a list of technological inventions over the 20th century by filtering Wikipedia titles, following [Asirvatham \(2024\)](#) and [Kalyani et al. \(2025\)](#) who showed their usefulness for building catalogs of technologies.⁵ Second, we use a Large Language Model (LLM) to summarize basic information about these inventions, including the time, location, and type of inventor (i.e., private company, individual inventor, government, or university/non-profit), with the prompt shown in Appendix Section [IA2.2](#). We also ask the LLM to report whether the company that was most successful in its initial implementation and commercialization was a young firm (less than ten years at the time) or an old incumbent, with the prompt shown in Appendix Section [IA2.3](#). As explained in the introduction, our work focuses on implementation, not simply the invention. Indeed, several well-known examples show that even when incumbents are able to invent new technologies, organizational rigidity can prevent them from effective implementation (e.g., Bell Labs invented transistors but did not commercialize them, Xerox developed personal computing technologies but did not commercialize them). Ultimately, new technologies need to be implemented to have an impact.

Figure 1 takes the technological inventions in the 20th century from the Wikipedia-based dataset, and plots the fraction that is most successfully implemented by young vs old firms. The first two bars restrict to technological inventions in the U.S., where the left bar (“All”) includes all inventions and the right bar (“Private”) includes only inventions by private companies. The LLM cannot clearly determine the answer or identify a commercializing firm for a larger share of inventions by individuals, universities, and governments (included in the “All” bar but not in the “Private” bar) compared to inventions by private companies, which is understandable. The second two bars show the corresponding results for all technological inventions in the world. Overall, we see a mix: a substantial amount of new technologies are successfully

⁵We download Wiki titles from [Wikimedia English Wikipedia Dumps](#) and use the steps in Appendix Section [IA2.1](#) to filter for technological inventions.

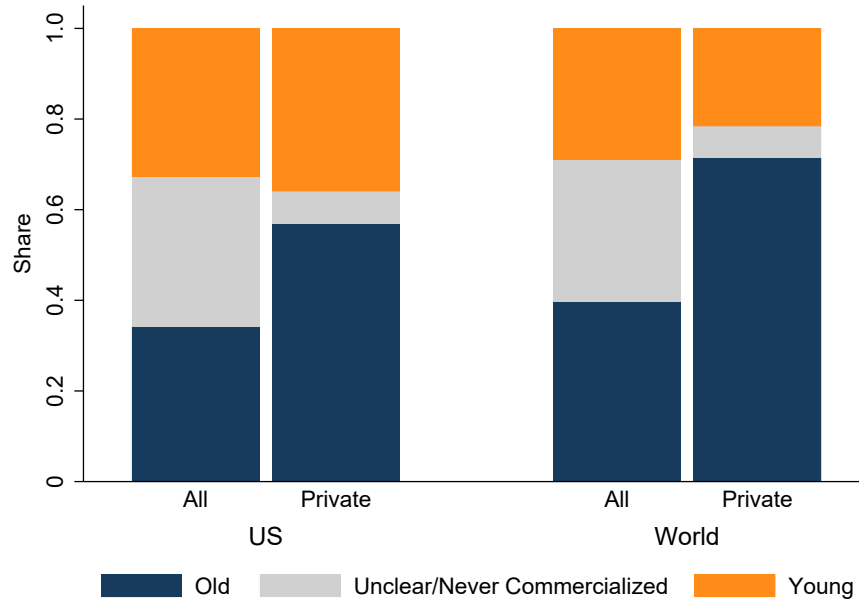


Figure 1. New Technologies Implemented by Young vs Old Firms (1900–2000)

Notes: This figure shows the fraction of technological inventions between 1900 and 2000 that are most successfully implemented by young firms (less than ten years) vs old firms (otherwise). The first two bars restrict to technological inventions in the U.S., where the left bar includes all inventions and the right bar includes only inventions by private companies. The second two bars show all technological inventions in the world, where the left bar includes all inventions and the right bar includes only inventions by private companies. The shares are normalized to one within each group. Colors correspond to firm age categories.

implemented by old firms, though some are not.

II.B Some Classic Examples

Business school classes have analyzed a number of cases featuring new technologies that incumbent firms succeeded or failed to implement. For example, Nylon (and polymers more generally) was invented and most successfully commercialized by DuPont, when the company was already more than 100 years old (Ndiaye, 2007). Antibiotics were most successfully commercialized by established companies such as Pfizer and Merck, which had existed for decades by that time. Television sets and jet engines were also produced most extensively by companies such as General Electric that had been founded long before their invention. More recently, Apple released the iPhone in 2007 when the company was over 30 years old.

Conversely, cars, aircraft, and semiconductors were most successfully produced by new companies at the time of their invention. Software was also mainly developed and commercialized by new companies, despite existing companies such as General Electric repeatedly trying to enter this domain (Cohan, 2022). Discount retail and discount airlines—which represent new business models though not necessarily

technological inventions—are also associated with new firms (O’Reilly and Tushman, 2021).

A common view in business school classes is that the former group largely represents new technologies that were compatible with existing organizational processes. Nylon had similarities with the production of rayon for DuPont; antibiotics benefited from large-scale fermentation that Pfizer and Merck already used; TVs and jet engines built on GE’s existing capabilities in manufacturing generators, turbines, and appliances; and iPhones follow from the production of other Apple devices. The latter group, however, led to organizational processes that differed from those in existing companies. The production of cars, aircraft, and semiconductors required a workflow that deviated from the procedures of their predecessors; software emphasized flexibility and maximizing upside, whereas hardware focused on attention to detail and minimizing downside; discount retailers and discount airlines relied on volume, whereas traditional retailers and airlines relied on margin.

How to capture the extent to which new technologies entail new organizational processes or priorities? These features are rich and difficult to measure in a uniform way. Our entry point is to extract information from the types of occupations required to implement new technologies. For example, as technological advancement leads to an increasing reliance on software in car manufacturing, the corresponding organizational processes will have to change given the different attributes of software vs hardware engineering that we later formalize as workstyle, which can be reflected by the shift in the occupation composition from hardware engineers to software engineers. Changes in the composition of occupations with different workstyles thus provide a window that can reveal corresponding changes in organizational processes. This angle has some limitations: subtle process changes may not be reflected in occupational composition, and the reliance on such data confines our main empirical analysis to the past two decades. However, this approach allows us to develop measurement methods that apply broadly across industries, with which we can perform systematic empirical tests.

III Model

We formalize the core hypothesis in a simple model, which helps guide our subsequent empirical analyses. The key ingredients of the model are as follows. First, each firm has two business models: a pre-existing business model (A), which is already in place, and a new business model (B), which the firm is seeking to adopt. Second, each business model employs workers in occupations $j \in \{1, \dots, J\}$. Each occupation has different productivity depending on the business model. Additionally, each occupation is characterized by a workstyle, ws_j , which does not impact its productivity but will govern organizational frictions between the two business models. Third, a given business model has a continuum of identical

operating units, each led by a manager who decides the occupation composition of the unit; we will explain why we need operating units later in Sections III.B and III.C. Finally, the firm owner oversees the choice of the number of operating units under each business model, while each manager supervises their own operating unit.

Our description of the model proceeds as follows. In Section III.A, we first describe outcomes for a firm consisting only of operating units under business model A , which we interpret as the status quo, before business model B is introduced. Taking these choices as given, in Section III.B we then describe the decision to expand into business model B when that opportunity arises, subject to organizational frictions between the old and new business models. The simple one-shot expansion assumption helps us illustrate clearly the key mechanisms of the model. In Section III.C, we provide a more general model with an explicit microfoundation for the organizational frictions, where the firm operates under business model A for some time, and the opportunity to adopt business model B arrives randomly. We explain how rules from the existing business model A build up over time as firms age, which codify institutional knowledge that facilitates its operations, but can increase the difficulty of adapting to new businesses that have different workstyles and require different approaches for solving problems. In Section III.D, we provide empirical support that rules increase with firm age using employee reviews from Revelio. Finally, in Section III.E, we use the simple model to derive explicitly the empirical predictions about the growth of young firms vs old firms that we test in subsequent analyses.

III.A Business Model A

We first describe the firm's optimization problem for the old business model A , before the emergence of the new business model B . For both, the resulting per-period optimization result remains valid in our later model extensions with dynamic features.

Technology and team composition Initially, firm i only operates units under business model A . Each unit is indexed by k , with one manager of business model A (henceforth, an A -manager) who is in charge of hiring workers among J occupations for that unit, as illustrated in Figure 2. More specifically, for each unit under business model A , the corresponding manager solves the following problem:

$$Q_{A,i,k} = \max_{\{l_{A,i,k,j}\}_{j=1}^J} \left(\sum_{j=1}^J \psi_{A,j}^{\frac{1}{\sigma}} \cdot (l_{A,i,k,j})^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}}, \quad (1)$$

$$s.t. \quad \sum_{j=1}^J l_{A,i,k,j} \leq \bar{l}. \quad [\lambda_{A,i,k}] \quad (2)$$

Here, $\{\psi_{A,j}\}$ captures the productivity vector across the J occupations, and \bar{l} captures the natural limit of team size that a manager can deal with.

The solution to problem (1) is as follows. For all k , we have $\lambda_{A,i,k} = \lambda_A = \left(\sum_{j=1}^J \psi_{A,j}\right)^{\frac{1}{\sigma-1}}$, and

$$l_{A,i,k,j} = \frac{\psi_{A,j}}{\sum_{j=1}^J \psi_{A,j}} \bar{l}. \quad (3)$$

This optimal composition of occupations implies the output per operating unit is

$$Q_{A,i,k} = Q_A = \left(\sum_{j=1}^J \psi_{A,j}\right)^{\frac{1}{\sigma-1}} \bar{l}. \quad (4)$$

Note that in each operating unit k across firm i in this industry, the manager will hire a team of workers with the same composition $\{l_{A,j}\}$ as in (3) and generate the same amount of output Q_A as in (4).

Workstyle of business model A We use ws_j to denote the workstyle associated with each occupation j . For firm i , the workstyle for its unit k under business model A is defined as

$$ws_A = ws_{A,i,k} \equiv \sum_{j=1}^J \frac{l_{A,i,k,j}}{l_{A,i,k}} ws_j, \quad (5)$$

where $l_{A,i,k} \equiv \sum_{j=1}^J l_{A,i,k,j}$. From (3) we see that $ws_{A,i,k} = ws_A$ is the same across all units k of firm i , and across all firms in the industry, simply because in each operating unit k its manager will hire a team of workers with the same composition $\{l_{A,j}\}$. The workstyle of business model A will play a role in shaping the organizational conflict when the new business model arrives.

Size of old business model A Because each unit is the same, at the firm level the problem is to choose how many operating units to have under business model A , taking the exogenous worker wage w and the A -manager's equilibrium wage (to be determined shortly) as given. We use $L_{A,i}$ to denote the number of operating units in business model A at firm i , which is also the total number of A -managers. The firm solves:

$$\max_{L_{A,i} \geq 0} \left(Q_A - w\bar{l} - w_A \right) L_{A,i} - \frac{L_{A,i}^2}{2\xi_i}. \quad (6)$$

The fixed effect of the firm ξ_i captures frictions that limit firm size besides organizational rigidities. It does not play a central role in our analysis, other than generating ex-ante heterogeneity in scale between firms under the old business model; all our main results hold if $\xi_i = \xi$ for all firms i in the industry. Let

$z_A \equiv Q_A - (w\bar{l} + w_A)$. The solution of (6) for the number of operating units of type A is

$$L_{A,i} = z_A \xi_i, \quad (7)$$

and the number of workers at the firm, denoted by $E_{A,i}$, is:

$$E_{A,i} = z_A \xi_i \bar{l}. \quad (8)$$

Equilibrium wage of A -managers We close the model by endogenizing the wage of A -managers, denoted by w_A . Assume that the total number of A -managers in the industry is exogenously given by L_A . The market clearing condition for A -managers with a unit mass of firms is $\int_0^1 L_{A,i} di = L_A$. Denote $\Xi \equiv \int_i \xi_i di$, then we have

$$w_A = (Q_A - w\bar{l}) - \frac{L_A}{\Xi}, \text{ and } z_A = \frac{L_A}{\Xi}.$$

Since L_A is fixed, the total employment of business model A in the industry will be fixed at $\bar{l}L_A$.

Finally, although we set up the problem of individual A -manager in (1) as a static one, the resulting static optimization is equivalent to a dynamic one in our setting since there are no intertemporal adjustment costs. Furthermore, this property continues to hold when we later introduce the arrival of business model B in Section III.B and the endogenous length of rules in Section III.C.⁶

III.B Business Model B

Later, a new business model B arrives, and firm i can now add new units under business model B , as illustrated in Figure 2.

Technology and team composition The new business is modeled as a different productivity vector $\{\psi_{B,j}\}$ over occupations. For simplicity, we assume that firm i cannot adjust existing business model A 's units alongside its decision to expand into business model B . In each operating unit k under business

⁶See Appendix Section IA6.1 for a formal proof. Implicitly, by adopting the two-layer structure where managers hire workers as in (1) and the firm owner hires managers as in (6), we rule out that the firm owner can dictate the worker composition in any individual operating unit (of type either A or B). Otherwise, organizational conflicts between models A and B will distort the choice of worker composition in a way that will depend on firm-specific frictions, making industry aggregation intractable.

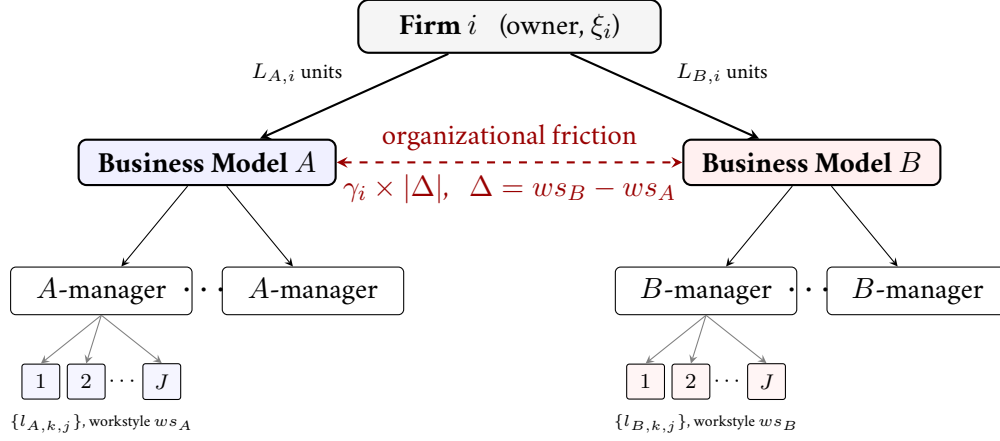


Figure 2. Model Structure

Notes: This figure illustrates the structure of firm i . Business model A is the pre-existing model; business model B is the new model. The firm owner chooses the number of operating units under each business model, $L_{A,i}$ and $L_{B,i}$, and ξ_i governs the overall scale of the firm. Each unit is led by a manager (an A -manager or B -manager) who hires workers across J occupations; $\{l_{A,k,j}\}$ and $\{l_{B,k,j}\}$ denote the employment of occupation j in unit k under each business model. The employment-weighted average workstyle under each business model is ws_A and ws_B , and $\Delta = ws_B - ws_A$ is the workstyle distance. The organizational frictions between the two models depend on the firm's rigidity parameter γ_i and $|\Delta|$.

model B , the manager (henceforth, a B -manager) solves a problem similar to that under business model A :

$$Y_{B,i,k} = \max_{\{l_{B,i,k,j}\}_{j=1}^J} \eta_i \cdot \underbrace{\left(\sum_{j=1}^J \psi_{B,j}^{\frac{1}{\sigma}} \cdot (l_{B,i,k,j})^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}}}_{Q_{B,i,k}} \quad (9)$$

$$s.t. \quad \sum_{j=1}^J l_{B,i,k,j} \leq \bar{l} \quad [\lambda_{B,i,k}] \quad (10)$$

The only difference from A 's problem in (1) is the presence of the term η_i , which captures the firm-level productivity of the business model B . As we will micro-found shortly, η_i reflects organizational frictions within the firm, the key economic mechanism in our setting.

Because each B -manager for operating unit k takes η_i as given, the solution structure is identical to that of business model A . We therefore have $\forall k, \lambda_{B,i,k} = \lambda_B = \eta_i \left(\sum_{j=1}^J \psi_{B,j} \right)^{\frac{1}{\sigma-1}}$, and

$$l_{B,i,k,j} = \frac{\psi_{B,j}}{\sum_{j=1}^J \psi_{B,j}} \bar{l}, \quad (11)$$

$$Q_{B,i,k} = Q_B = \left(\sum_{j=1}^J \psi_{B,j} \right)^{\frac{1}{\sigma-1}} \bar{l}, \quad (12)$$

$$Y_{B,i,k} = Y_{B,i} = \eta_i Q_B. \quad (13)$$

Workstyle difference and productivity η_i of business model B For each operating unit k of type B , we define the workstyle difference with business model A as:

$$\Delta_{i,k} \equiv \sum_{j=1}^J \frac{l_{B,i,k,j}}{l_{B,i,k}} w_{S_j} - w_{S_A}, \quad (14)$$

where $l_{B,i,k} \equiv \sum_{j=1}^J l_{B,i,k,j}$ and w_{S_A} is given by (5). Here, the first term is the employment-weighted average workstyle in each new unit k of business model B . One can rewrite (14) as:⁷

$$\Delta_{i,k} = \sum_{j=1}^J \frac{l_{B,i,k,j}}{l_{B,i,k}} \delta_j, \quad \text{with} \quad \delta_j \equiv w_{S_j} - w_{S_A}.$$

As before, we use $L_{B,i}$ to denote the total number of operating units under business model B (to be solved shortly). We can then define the average distance in workstyle between existing business model A and the new business model B :

$$|\Delta|_i \equiv \frac{1}{L_{B,i}} \int_0^{L_{B,i}} |\Delta_{i,k}| dk. \quad (15)$$

In general, the productivity η_i of business model B in (9) may depend either on collaboration with the existing business model A (e.g., sharing know-how or functional units)—which is the focus of our main model; or on support from headquarters (e.g., providing state-of-the-art facilities and infrastructure)—which we discuss in Appendix Section IA6.2. We assume that η_i takes the following form:

$$\eta_i \equiv \eta - \gamma_i |\Delta|_i. \quad (16)$$

Here, $\eta > 1$ is a parameter that reflects potential efficiency gains when the business model B is developed inside the incumbent firm. The second term of (16) reflects the cost from organizational frictions, which is our key focus. It is greater when the workstyle distance between old and new business models $|\Delta|_i$ is larger, in which case organizational conflicts are more likely to arise. The sensitivity of this cost to the workstyle distance depends on the parameter $\gamma_i > 0$, which captures the severity of frictions in firm i for dealing with organizational conflicts. We provide microfoundations for this specification in Section III.C.

Using the solution to the problem of each operating unit described above, especially (11), we observe that the average workstyle difference is independent of firm i :

$$\forall k, \quad \Delta_{i,k} = \sum_{j=1}^J \frac{\psi_{B,j}}{\sum_{j'=1}^J \psi_{B,j'}} \delta_j \equiv \Delta. \quad (17)$$

⁷Note that so long as $w_{S_j} \neq w_{S_{j'}}$ for at least one pair (j, j') , then at least one of the δ_j 's is strictly negative.

Thus $\Delta_{i,k} = \Delta$ is constant across firms (in the particular industry we are studying) and across operating units of type B within a firm. Then, from (16), this observation implies that the equilibrium firm-level value of η_i is:

$$\eta_i = \eta - \gamma_i |\Delta|. \quad (18)$$

As explained right after (16), every B -manager takes η_i as given when solving problem (9). As evident from (11), the upshot of our assumption that individual managers of operating units take η_i as given is that the resulting optimal team composition in each B unit is independent of the firm-specific organizational friction parameter γ_i . That is, under our microfoundation provided in Section III.C, the productivity of business model B , $\eta_i = \eta - \gamma_i |\Delta|$ —especially $|\Delta|$ as in (15)—is independent of $\Delta_{i,k}$, which is the occupation composition of the individual unit k that manager k controls. Instead, it depends on the overall occupation composition of B . Of course, organizational frictions will affect the number of units under business model B (hence the number of B -managers) for each firm, as shown in (19). We make this assumption for analytical convenience only, as homogeneous composition across all firms within an industry allows us to derive industry-level predictions and connect to empirical analyses in a clean way.

Size of new business model B The owner of firm i decides the optimal size of business model B by choosing the number of B -managers to solve:⁸

$$\max_{L_{B,i} \geq 0} \left(\eta_i Q_B - w\bar{l} - w_B \right) L_{B,i} - \frac{L_{B,i}^2}{2\xi_i}. \quad (19)$$

Here, ξ_i is the same quadratic cost of employing managers as in (6) for business model A .

Given the B -manager's equilibrium wage w_B and productivity $\eta_i Q_B$ in (13), we denote $z_{B,i} \equiv (\eta - \gamma_i |\Delta|) Q_B - (w\bar{l} + w_B)$. The solution for the number of operating units of type B , $L_{B,i}$, and the number of total workers under business model B , $E_{B,i}$, is:

$$L_{B,i} = z_{B,i} \xi_i, \quad \text{and} \quad E_{B,i} = z_{B,i} \xi_i \bar{l}. \quad (20)$$

Equilibrium wage of B -managers Finally, as before, we can close the model to pin down the equilibrium wage of B -managers, w_B , by assuming that there is a fixed number of B -managers. Define:

$$\Gamma \equiv \frac{\int_i \xi_i \gamma_i di}{\int_i \xi_i di}.$$

⁸Implicitly we assume that the firm owner will not adjust the occupation composition of business model A , and neither can contract with each B -manager on the occupation composition of each operating unit.

Solving for w_B , we obtain:

$$w_B = (\eta - \Gamma |\Delta|) Q_B - w\bar{l} - \frac{L_B}{\Xi}, \quad (21)$$

$$z_{B,i} = (\Gamma - \gamma_i) |\Delta| Q_B + \frac{L_B}{\Xi}. \quad (22)$$

Equation (22) gives one of the key properties of our model: the equilibrium size of new business model B is firm-dependent and decreases with the severity of firm i 's organizational frictions.

III.C A Microfoundation of Organizational Frictions between Business Models

In Section III.B, we assumed that the productivity of B -managers is affected by the workstyle differences between new and existing business models, via a simple functional form $\eta_i = \eta - \gamma_i |\Delta|$. Crucially, the B -manager efficiency depends on the difficulty of resolving organizational conflict in firm i , which is indexed by the parameter γ_i . We now provide a microfoundation for this assumption, and illustrate how γ_i can increase with firm age.

The key mechanism rests on the idea that the old business model A gradually accumulates rules over time, which help solve the types of problems typically encountered by A -workers. However, these rules may not be suitable for workers in business model B , especially when workstyles differ across the two business models. As a result, disagreements emerge among workers operating under the two business models and are escalated to their managers. Resolving these disagreements requires time and managerial effort, effectively reducing the productivity of business model B . Moreover, when the body of rules accumulated under the existing business model A are longer and more complex, the conflict resolution process becomes more time consuming for business model B , leading to greater inefficiency.

Rules for business model A We consider a simple dynamic model before the emergence of business model B . Time is discrete and firm owners discount the future at a constant rate ρ .⁹

Firm i initially operates only under business model A . In each period, with constant probability $p \in (0, 1)$, business model B is expected to arrive. Before business model B has arrived, the firm functions the same way as in Section III.A. In particular, the occupation composition in each operating unit is decided by managers who do not take into account externalities that the composition of their unit might cause on the rest of the firm.

Each period before business model B arrives, the firm may run into a new incident $d_{i,t} \in \{1, 2, \dots, R\}$,

⁹See Appendix Section IA6.1 for a formal statement of the dynamic model.

where R is the total number of potential incident types. Each type of incident occurs with equal probability $1/R$, and incidents are i.i.d. over time. Suppose that $r_{i,t} \leq R$ distinct incidents from the list $\{1, 2, \dots, R\}$ have occurred to firm i up to the current period t . There are two possibilities.

- With probability $1 - p$, the firm continues to operate only units under business model A . An incident then arises. One of the following two events could potentially occur:
 - With probability $\frac{r_{i,t}}{R}$, the incident that arises is an old one, and the firm can resolve it with existing rules at no cost;
 - With probability $1 - \frac{r_{i,t}}{R}$, the incident that arises is a new one. The firm needs to pay a cost c to resolve it, and adds new rules so that in the future this incident can be resolved at no cost, by following these new rules.
- With probability p , business model B arrives. For simplicity, we assume that there are no incidents in business model A going forward, but that incidents in business model B start to occur.

We assume that the list of rules is managed by A -managers. This is consistent with the premise throughout our model that managers are in charge of operational decisions, including worker occupation composition as in Section III.A and III.B. Because adding rules always helps the efficiency of the business model in question, A -managers will keep appending the rules following new incidents. That is, the rule length $r_{i,t}$ grows monotonically over time until $r_{i,t} = R$ so that set of rules is complete.¹⁰ In our setting the accumulation of rules helps formalize solutions to problems encountered by business model A and build its organizational capital (Rajan, 2012; Levitt, List, and Syverson, 2013). However, as we model next, these rules can hurt the performance of business model B , in that the length $r_{i,t}$ of the rule list for business model A can negatively affect the productivity of business model B .

Figure 3 summarizes the timing within each period before business model B arrives.

Frictions between old and new business models Now suppose that business model B has arrived, and that firm i has $r_{i,t}$ rules developed by the existing business model A . Moreover, because there are no new incidents to A going forward, the length of the list of rules becomes fixed; we denote it by $r_i = r_{i,t}$, where t is the time of first arrival of business model B .

The choice of worker composition in each operating unit of business model B is the same as described in the simple one-shot model. In particular, each B -manager takes the productivity parameter η_i as given,

¹⁰In an earlier version we considered the possibility that the rule list is managed by the firm owner. In that case, the owner would take into account the negative effect of rule length on the productivity of business model B , and hence the rule length would stop growing once $r_{i,t}$ hits some endogenous level R^* . This does not change the key property that older firms tend to have longer rules.

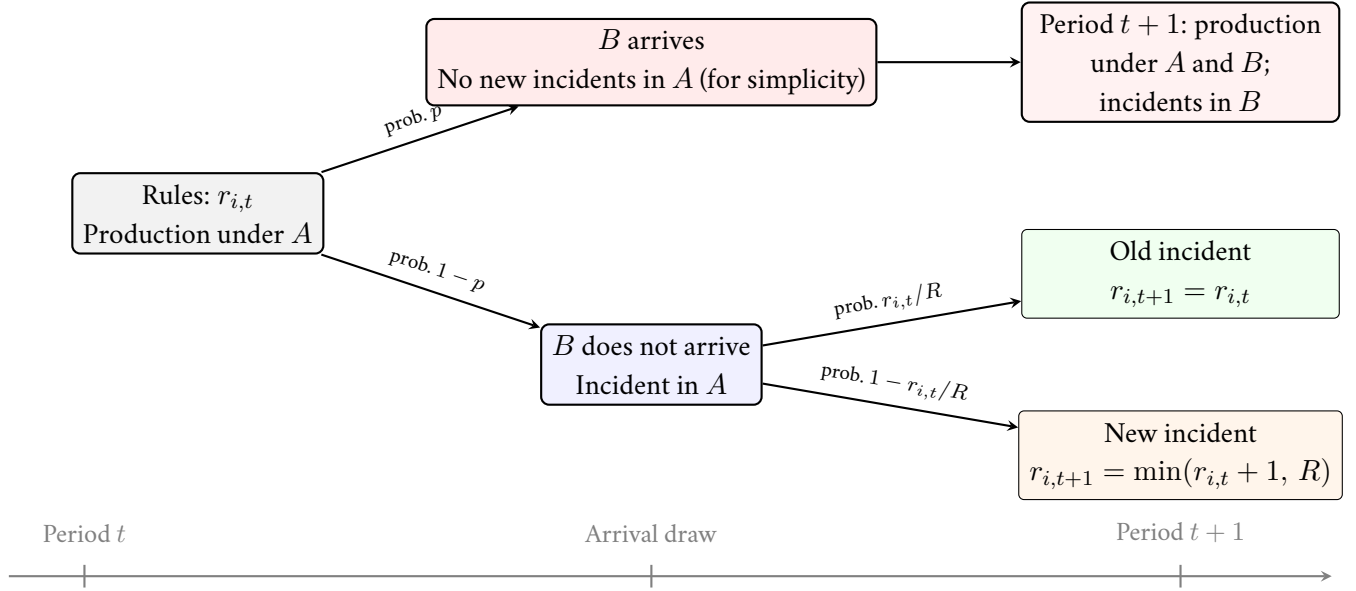


Figure 3. Timeline of the Dynamic Model

Notes: This figure illustrates the timing of events in the dynamic model.

resulting in a workstyle difference $\Delta_{i,k} = \Delta$ with respect to business model A that is identical across business model B units.

We follow [Garicano \(2000\)](#) in modeling the business model B 's productivity η_i as a function of the B -manager's effective time. Each B -manager is endowed with one unit of time, which is used to address workers' normal inquiries, assign them to work activities, and monitor their performance ([Bloom and Van Reenen, 2007](#); [Minni, 2026](#)); this shapes the team's baseline productivity. But occasionally "incidents" may occur, which affects the effective management time and hence the team's productivity η_i .

For each operating unit k in business model B , new incidents may arise with probability $z \in (0, 1)$, and workers in business model B (B -workers) propose a solution. Workers in business model A (A -workers) disagree (agree) with the solution with probability $|\Delta|(1 - |\Delta|)$. The key is that the probability of disagreement is increasing in the distance between the workstyles of A and B ; in words, if the workstyles of A and B misalign, they are more likely to disagree.¹¹

When workers agree, A -workers use their know-how to help B -workers resolve incidents more efficiently; we think of these incidents as problems that arise at the level of operating facilities, and the knowledge sharing captures efficiency gains in the spirit of [Levitt, List, and Syverson \(2013\)](#). We assume that this reduces the B -manager's involvement and saves managerial time by $\hat{\eta}$, which is a positive constant.¹²

¹¹Recall that $\Delta_{i,k} = \Delta$ for all firm i and operation unit k . For ease of exposition we assume that $|\Delta| \in (0, 1)$. If $|\Delta| > 1$ then any scaled version of $|\Delta|$ as the probability of conflict will work.

¹²When workers in two business models A and B agree, these incidents tend to be more routine; therefore one may posit

As a result, in expectation, these synergies improve business model B 's productivity by $z(1 - |\Delta|)\hat{\eta}$. These synergies can also come from sharing functions such as IT, human resources, or branding and marketing.

When workers disagree—which we call a “conflict”—they cannot resolve the issue on the spot and escalate it to their managers. All B -managers participate in the resolution process, as these conflicts provide learning opportunities across all units under business model B . They form committees to show that the proposed solutions from B -workers are justified relative to the existing rules established under business model A . Although existing rules may not be useful for the problem at hand, A -managers believe in them and insist that these rules should be followed. B -managers must therefore expend effort persuading them that the current situation warrants a deviation from established rules. The longer and more complex the rules accumulated under business model A , the more effort B -managers must devote to conflict resolution.

We assume that each committee requires $\hat{\gamma}(r_i)$ units of time, with $\hat{\gamma}(\cdot)$ a strictly increasing function. Recall that firm i has $L_{B,i}$ units under business model B . Because each conflict is i.i.d., the total number of conflicts is $\int_0^{L_{B,i}} z |\Delta| dk = z |\Delta| L_{B,i}$, and hence the total conflict resolution time for B -managers is $z L_{B,i} |\Delta| \hat{\gamma}(r_i)$.¹³ Because the committee work is shared among all B -managers with a total measure of $L_{B,i}$, each individual B -manager spends $z \hat{\gamma}(r_i) |\Delta|$ units of time in resolving the conflict. Therefore the effective management time of B -manager, taking into account the benefit from agreement occurs and the cost from disagreement, is

$$\eta_i = 1 + \underbrace{z(1 - |\Delta|)\hat{\eta}}_{\text{benefit of agreement}} - \underbrace{\frac{1}{L_{B,i}} \int_0^{L_{B,i}} z |\Delta| \hat{\gamma}(r_i) dk}_{\text{cost of disagreement}} = \underbrace{1 + z\hat{\eta}}_{\eta} - \underbrace{z(\hat{\gamma}(r_i) + \hat{\eta})}_{\gamma(r_i)} |\Delta|. \quad (23)$$

That is, taking into account the organizational conflict, the productivity of each operating unit in business model B becomes $\eta - \gamma(r_i) |\Delta|$. By rewriting $\gamma(r_i)$ as γ_i , this is exactly our assumption in (18).

The organization cost in (23) affects per-period profit for business model B given in (19).¹⁴ In short, the microfoundation of organizational frictions in (23) says that the accumulation of rules from the existing business models exacerbates conflicts stemming from workstyle differences between old and new business models. Although rules help implement existing routines and address familiar problems, they can increase the difficulty of adapting to new businesses that require different approaches for solving problems. Since rules build up over time, older firms carry a larger volume of them—and accordingly

that $\hat{\eta}$ is increasing in the length of rules. The same model mechanism goes through if we assume that $\hat{\eta}$ is monotonically increasing in r_i .

¹³Technically, the probability of incidents for each operating unit in business model B is $z \cdot dk$.

¹⁴Once an incident in business model B (whether it is a conflict or not) is resolved, the resulting solution is incorporated into business model B 's rule set for daily operations, analogous to the procedure followed under business model A . Because there is no further arrival of new business models, it is always optimal to add to the rules for business model B . And, for simplicity, we assume $R \rightarrow \infty$ so that the length of the rule list in business model B does not matter for the value.

face greater inefficiency and delay in resolving conflicts, especially those arising from new initiatives that challenge established practices. Older firms, often weighed down by outdated rules and entrenched norms, are more likely to exhibit higher γ_i .¹⁵ Accordingly, old firms face greater challenges in implementing the new business model when the associated workstyle change is large, but not as much when the workstyle change is small.

Can business model B be insulated from business model A ? In the microfoundation associated with (23), workers from business models A and B interact with each other, and these interactions have costs and benefits. If the costs of disagreement and its slow resolution are too severe relative to the benefits of working together, a natural question is whether firms can avoid conflicts by separating the operation of the new business model B entirely from that of business model A .

In Appendix Section IA6.2 we consider the possibility of establishing a distinct division for business model B . Although workers from the two divisions no longer interact directly with each other, following the tradition of the incomplete contracting literature we assume that both divisions (i.e., division managers) report to the same headquarters, which retains the authority to allocate certain firm-level resources, such as IT infrastructure or discretionary budgets. We show that conflicts could still arise between the two divisions and generate an agency cost similar to (18), based on the mechanism of “contract as reference point” in Hart and Moore (2008). In a nutshell, larger workstyle differences can lead to more grievances from division A ,¹⁶ triggering division A ’s sabotage of division B through headquarters (e.g., by reporting that division B has engaged in misbehavior and requesting an investigation). In older firms, the existing division is more experienced and influential, leading to larger efficiency losses for division B . These frictions across divisions apply as long as the divisions are run by the same headquarters, without selling off B as an independent firm (if firm i spins off B as another firm, then its growth no longer contributes to the employment, sales, or valuation of firm i that we analyze in the data and does not count as firm i implementing B).

The organizational frictions we analyze may also shed light on why financial investors like private equity firms can manage portfolio companies with distinct styles, while doing so is difficult for operating

¹⁵Our microfoundation models organizational friction as wasteful “committee time” for conflict resolution, but γ_i can be interpreted more broadly. A high γ_i may reflect rigid organizational structures—hierarchical approval processes, siloed departments, and little tolerance for bottom-up experimentation—that raise the cost of resolving disagreements between old and new business models. For instance, proposals that challenge existing practices must pass through multiple approval layers, each controlled by managers invested in preserving the status quo. Such rigidity not only slows decision-making but may also dilute or suppress new ideas before they reach implementation.

¹⁶Given the organization cost, business model B with larger workstyle difference is endogenously more profitable, leading to less favorable resource allocation toward division A . Following the logic of Hart and Moore (2008), this ex post shortfall (relative to the maximum possible allocation which is one) causes aggrievement from division A managers. We also discuss other channels through which larger workstyle differences lead to more aggrievement from division A in Appendix Section IA6.2.

companies. Indeed, there are likely good reasons that companies like X, Tesla, and SpaceX are designed to be separate entities rather than a combined operating company. When these companies are operationally and financially separate, they can avoid internal organizational conflicts due to style differences that we model in the above, and avoid conflicts due to sharing resources in Appendix Section [IA6.2](#).

III.D Firm Age and Rules in the Data

The microfoundation in Section [III.C](#) associates older firms with more rules. To support this observation, we collect information through employee reviews from Revelio ([Revelio Labs, 2025](#)). First, we use LLM to screen for reviews that discuss rules at the employer company. Second, we use LLM to label the relevant reviews as those that mention the presence of rules vs those that mention the lack of rules; Appendix Section [IA3.1](#) shows our prompt. For each firm (represented by GVKEY) in Compustat, we estimate its age following the procedure in Appendix Section [IA5](#) (which is also used in the empirical analyses in Sections [V](#) and [VI](#)), and we calculate the intensity of rules in a GVKEY-year as:¹⁷

$$\text{Rule Index} = \frac{(\# \text{ of reviews mentioning the presence of rules} - \# \text{ of reviews mentioning the lack of rules})}{\# \text{ of total reviews}}. \quad (24)$$

Similarly, we use employee reviews to support the observation that old firms have more meetings, committees, and layers of approval that take time. We detail our prompt in Appendix Section [IA3.2](#). First, we use LLM to screen for reviews that discuss meetings, committees, and layers of approval at the employer company. Second, we use LLM to label the relevant reviews as those complaining the company has too many such activities that waste time vs praising the company for being efficient and not having too many such activities. We calculate the intensity of meetings that waste time in a GVKEY-year as:

$$\text{Meeting Index} = \frac{(\# \text{ of reviews complaining too many} - \# \text{ of reviews praising not too many})}{\# \text{ of total reviews}}. \quad (25)$$

Figure [4](#) shows binscatter plots of the relationship between the Rule Index and the Meeting Index with respect to age. We see that both are strongly increasing with firm age.

¹⁷We map Revelio company ID RCID to GVKEY in a given year using parent-subsiary bridge provided by Revelio to us directly.




Figure 4. Firm Age, Rules, and Meetings that Waste Time

Notes: This figure shows binscatter plots of the Rule Index in (24) in Panel A and the Meeting Index in (25) in Panel B in 10 equal-sized bins based on log firm age for Compustat firms. The Rule Index and Meeting Index are constructed from employee reviews collected by Revelio. We absorb year fixed effects.

III.E Connection to Empirical Analyses on Firm Growth

We now derive the testable implications of the model for the growth of young and old firms when new technologies induce changes in workstyle. We focus on the simple, one-shot model (Sections III.A and III.B) and consider firms in multiple industries indexed by n , where the industries differ only in their occupation productivity vectors $\{\psi_{A,j,n}, \psi_{B,j,n}\}_{j=1}^J$.

Firm growth and change in workstyle The following result summarizes the growth rate of firms, which is the basis of our empirical predictions.

Result 1. *In each industry n , the equilibrium growth rate of firm i is given by:*

$$g_{i,n} = \mu_n - \lambda Q_{B,n} |\Delta|_n \gamma_i, \quad (26)$$

where $\lambda \equiv \Xi/L_A$ and $\mu_n \equiv L_B/L_A + \lambda Q_{B,n} |\Delta|_n \Gamma$ are industry-wide constants, and:

$$|\Delta|_n \equiv \left| \sum_j \left(\frac{\psi_{B,j,n}}{\sum_{j'} \psi_{B,j',n}} - \frac{\psi_{A,j,n}}{\sum_{j'} \psi_{A,j',n}} \right) w_{S_j} \right| \quad (27)$$

is the workstyle distance between business models A and B in industry n . The change in the industry-level employment-weighted average workstyle is proportional to this distance:

$$|d\bar{w}_S|_n = \frac{L_B}{L_A + L_B} |\Delta|_n. \quad (28)$$

Proof. See Appendix Section [IA6.3.1](#). □

This result indicates that firms with higher γ_i grow more slowly when the workstyle change in the industry $|d\bar{w}s|_n$ is larger. We now express this observation as a reduced-form empirical regression.

Empirical regression We build on the previous result about the relationship between firm growth and workstyle change, and analyze the following regression in the data:

$$g_{i,n} = \nu + \zeta (|d\bar{w}s|_n \times \text{age}_i) + \kappa |d\bar{w}s|_n + \rho \text{age}_i + \varepsilon_{i,n}, \quad (29)$$

where $g_{i,n}$ is a measure of firm-level growth (of employment, in the model), $|d\bar{w}s|_n$ is the change in industry n 's weighted average workstyle, and age_i is the age of firm i .¹⁸ The main coefficient of interest is ζ . To help interpret the regression coefficients, we make the following assumption.

Assumption 1. *The joint distribution of (γ_i, age_i) is identical across industries n , and is independent of all industry-level variables (including $Q_{B,n}$, $|\Delta|_n$, and $|d\bar{w}s|_n$).*

The microfoundation in Section [III.C](#) emphasizes that organizational frictions increase with age, which is also our focus here and in the remainder of the paper. However, our derivations also apply to other exogenous firm characteristics that affect γ_i .

Result 2. *Under Assumption 1, the regression coefficient $\hat{\zeta}$ in Equation (29) is:*

$$\hat{\zeta} = -\Phi \times \frac{\text{cov}(\gamma_i, \text{age}_i)}{\text{var}(\text{age}_i)}, \quad \text{where} \quad \Phi \equiv \frac{\Xi}{L_A} \left(1 + \frac{L_A}{L_B}\right) \times \frac{\text{cov}(Q_{B,n} | \Delta|_n, |\Delta|_n)}{\text{var}(|\Delta|_n)}. \quad (30)$$

Proof. See Appendix Section [IA6.3.2](#). □

The economic content of Result 2 is as follows. The coefficient $\hat{\zeta}$ is the product of two terms: the industry-level scalar Φ , which captures the strength of the relationship between workstyle change and the business model B 's productivity; and the covariance between organizational frictions and firm age. When $\Phi > 0$, the estimated coefficient $\hat{\zeta}$ is negative if and only if older firms have larger organizational frictions. Under the condition $\Phi > 0$, this regression provides a direct test of our core hypothesis that occupational misalignment between different business models hampers growth for organizationally rigid firms.

There are good economic reasons to think that $\Phi > 0$ is the empirically relevant case. The scalar Φ is positive whenever $\text{cov}(Q_{B,n} | \Delta|_n, |\Delta|_n) > 0$, i.e., whenever industries with larger workstyle changes

¹⁸We focus on a single cross-sectional regression with growth rates computed from before to after the arrival of business model B . In the empirical analyses in Section [V](#), we will use firm growth and industry workstyle change over five years.

also tend to have more productive new business models. Intuitively, if firms are undertaking business model B , it is because the new business is sufficiently productive to be worth the cost of the organizational cost of the workstyle adjustment.¹⁹

In Appendix Section IA6.3.4, we provide a generalization of the regression in (29) to a multivariate specification that includes additional industry-level and firm-level controls. The coefficient of interest ζ retains the same multiplicative structure (Result IA5), with the bivariate regression coefficients replaced by their multivariate counterparts that partial out the additional controls. This generalization does not change the economic interpretation of the interaction coefficient between predicted workstyle change and firm age described in Result 2.

Who implements business model B ? In Result 1, we focus on the growth *rates* of old and young firms, which we test empirically in the following. A complementary question, motivated by the stylized facts in Section II, is about *levels*: under what circumstances will most of the implementation of business model B be done by young firms? One can show (see Appendix Section IA6.3.5, Result IA6) that the share of total business model B output produced by firm i in industry n is

$$\sigma_{B,i} = \frac{\xi_i}{\Xi} \omega(\gamma_i), \quad \omega(\gamma_i) \equiv \frac{\eta_i z_{B,i}}{\bar{D}}, \quad (31)$$

where \bar{D} is an industry-level constant. The function $\omega(\gamma_i)$ is strictly decreasing: firms with greater organizational frictions produce a smaller share of B output, both because their B -managers are less efficient (lower η_i) and because they operate fewer B -units (lower $z_{B,i}$). This effect is amplified when the workstyle distance $|\Delta|$ is large, since both η_i and $z_{B,i}$ become more sensitive to γ_i . Thus, when new technologies in an industry generate large changes in workstyle ($|\Delta|$), more will be implemented by firms with low organizational frictions, i.e., young firms which do not have a long list of rules that impede adaptation. Conversely, when new technologies do not shift workstyle by much, organizational frictions matter less and implementation is spread more evenly across firms regardless of age. This is consistent with the evidence in Section II, which documents a mix of young and old implementers across technologies.

IV Data and Measurement

We aim to measure industry-level workstyle changes $|d\bar{w}s|_n$ induced by new technologies. To do so, we combine two sets of data. First, we obtain workstyle at the occupation level using the O*NET “Work

¹⁹Appendix Section IA6.3.3 provides a formal sufficient condition: when the scale of the productivity improvement ($Q_{B,n}$) and the magnitude of the workstyle change ($|\Delta|_n$) are independent across industries, $\Phi > 0$ holds automatically.

Styles” dataset, which records the importance of 16 characteristics for the occupation (e.g., attention to detail, self control, stress tolerance, adaptability/flexibility, innovation). Second, we follow [Kogan et al. \(2024\)](#) to predict changes in occupation composition in each industry due to new technologies (based on patents). In a nutshell, employment of an occupation in an industry decreases (increases) when the occupation’s routine (nonroutine) tasks are similar to new technologies in the industry. Finally, we combine these two steps: new technologies predict future employment composition of different occupations in an industry, which then induces changes in the overall (occupation employment weighted average) workstyle in the industry. Our approach using information from occupation composition builds on prior work linking technology to changes in occupation composition ([Bresnahan, Brynjolfsson, and Hitt, 2002](#); [Autor, Levy, and Murnane, 2003](#); [Acemoglu and Autor, 2011](#)), and recent research showing that similarity in occupations across industries relates to their compatibility in the case of firms’ horizontal expansions and mergers ([Lee, Mauer, and Xu, 2018](#); [Beaumont, Hebert, and Lyonnet, 2025](#)).

IV.A Measuring Workstyle

We first describe the measurement of workstyle at the occupation level using O*NET data by the U.S. Department of Labor. The O*NET dataset describes the key characteristics of occupations. The information comes from surveys of workers and occupational experts. The occupations in O*NET use a slight variant of the Bureau of Labor Statistics’ Standard Occupational Classification (SOC) codes.

Workstyle for occupation j For each occupation, the O*NET Work Styles module provides the importance score of 16 characteristics, on a scale of 1 to 5. The characteristics are achievement/effort, persistence, initiative, leadership, cooperation, concern for others, social orientation, self-control, stress tolerance, adaptability/flexibility, dependability, attention to detail, integrity, independence, innovation, and analytical thinking. We use \mathbf{ws}_j to denote this 16×1 vector at the occupation level. We download the data from the O*NET website ([U.S. Department of Labor, Employment and Training Administration, 2025](#)).²⁰

Workstyle in industry n and year t We construct the industry-level workstyle in each year t as:

$$\mathbf{WS}_{n,t} = \sum_{j=1}^J \frac{l_{j,n,t}}{l_{n,t}} \mathbf{ws}_j. \quad (32)$$

²⁰For each occupation and characteristic, we take the average value across unique survey waves if there are multiple. If we only use the first value, then some occupations will have missing values in the early years.

We use the Occupational Employment and Wage Statistics (OEWS) dataset by U.S. Bureau of Labor Statistics (BLS) to measure the employment of occupation j in industry n and year t , $l_{j,n,t}$, as well as the industry’s total employment, $l_{n,t} = \sum_j l_{j,n,t}$. We use data at the three-digit NAICS level, which are consistently available since 2003. We download the data from BLS website ([U.S. Bureau of Labor Statistics, 2025](#)).

Workstyle change in industry n from year t to year $t + 5$ We define workstyle change in industry n from year t to year $t + h$ as the Euclidean distance between $\mathbf{WS}_{n,t}$ and $\mathbf{WS}_{n,t+h}$:

$$|d\overline{ws}|_{n,t,h} \equiv \|\mathbf{WS}_{n,t+h} - \mathbf{WS}_{n,t}\|, \quad (33)$$

which is the empirical counterpart to $|d\overline{ws}|_n$ in the model. We use $h = 5$ in our baseline analyses. We standardize $|d\overline{ws}|_{n,t,h}$ to unit variance and zero median to facilitate the assessment of economic magnitude in the regressions.

IV.B Technology and Industry Workstyle Change

We then connect $l_{j,n,t+5}$ and correspondingly $\mathbf{WS}_{n,t+5}$ to new technologies building on the methodology of [Kogan et al. \(2024\)](#), based on breakthrough patents ([Kelly et al., 2021](#)). The core idea is to predict future employment by occupation in an industry using the similarity between occupation tasks (from task description text) and new technologies (from patent text); Appendix Section [IA4](#) provides further details on the similarity computation. We use 183,213 U.S. breakthrough patents granted between 2003 and 2016 (the end of breakthrough patent data).

Calculate patents’ similarity with occupation j ’s routine & nonroutine tasks We obtain the tasks of each occupation according to the O*NET-SOC 2000 taxonomy ([U.S. Department of Labor, Employment and Training Administration, 2005](#)), so that the task descriptions predate the sample period of our analyses. We then query LLM to label them as routine tasks vs non-routine tasks.²¹ For each occupation, we combine the descriptions of routine tasks and nonroutine tasks, and will use each set to calculate the similarity with patent text later. We also compute the share of tasks that are routine or non-routine for each occupation.

²¹We follow [Kogan et al. \(2024\)](#) in querying LLM (GPT-4o) for this classification. The prompt is: “A routine task can be defined as follows: A routine task involves carrying out a limited and well-defined set of work activities, those that can be accomplished by following explicit rules. These tasks require methodical repetition of an unwavering procedure, and they can be exhaustively specified with programmed instructions and performed by machines. Tell me whether the following task is primarily routine or primarily non-routine; and explain your reasoning in one sentence.

Task: *task statement text from O*NET*

Output your answer in JSON like the following format: {"answer": "primarily routine/primarily non-routine", "reasoning": "your reasoning"}

The median occupation has 16 tasks and 33.3% are routine.²²

We follow [Kogan et al. \(2024\)](#) to represent each document—including each patent and each occupation’s combined routine (nonroutine) task descriptions— X_i as a weighted average of its word embeddings x_k :

$$X_i = \sum_k q_{i,k} x_k, \quad (34)$$

where $q_{i,k}$ is the Term Frequency-Inverse Document Frequency (TF-IDF) weight,²³ and word embeddings x_k are obtained using the *GloVe* model.²⁴ Similar to [Kogan et al. \(2024\)](#), we pre-clean the patent and occupation text to keep only verbs and nouns and remove stop words before calculating the embeddings, as explained in detail in Appendix Section [IA4.1](#).

We then calculate the cosine similarity $s_{p,j}^r$ between a patent p and the routine or non-routine component of occupation j :

$$s_{p,j}^r = \frac{\mathcal{X}_p}{\|\mathcal{X}_p\|} \cdot \frac{\mathcal{X}_j^r}{\|\mathcal{X}_j^r\|}, \quad r \in \{R, NR\}. \quad (35)$$

We follow [Kogan et al. \(2024\)](#) to remove year fixed effects (for routine and non-routine metrics and for each year separately), and impose sparsity on $s_{p,j}^r$ by setting values below the 80th percentile to zero (then rescaling the values to $[0, 1]$) to obtain an adjusted similarity measure $\tilde{s}_{p,j}^r$ (see Appendix Section [IA4.1](#) for details).

Sum over patents in industry n and year t to get occupation j ’s exposure $\xi_{j,n,t}^R$ and $\xi_{j,n,t}^{NR}$ We sum over similarity between occupation j and all granted patents assigned to industry n in year t according to the mapping between Cooperative Patent Classification (CPC) and three-digit NAICS industries provided by [Goldschlag, Lybbert, and Zolas \(2016\)](#). We refer to the total similarity between occupation j ’s routine (nonroutine) tasks and patents in industry n year t as $\xi_{j,n,t}^R$ ($\xi_{j,n,t}^{NR}$). Specifically, the exposure of occupation j in industry n to technology at year t is calculated as:

$$\xi_{j,n,t}^r = \theta_j^r \log \left(1 + \sum_{p \in \mathcal{P}_t} \pi_{p,n} \tilde{s}_{p,j}^r \right), \quad r \in \{R, NR\}, \quad (36)$$

²²The 2000 O*NET classification contains 15,643 different tasks, 98.3% of which are specific to a particular occupation in the O*NET-SOC classification. Examples of tasks include: "Manage and treat common health problems, such as infections, influenza and pneumonia, as well as serious, chronic, and complex illnesses, in adolescents, adults, and the elderly"; "Measures and marks location of studs, leaders, and receptacle openings, using tape measure, template, and marker"; "Receive mortgage, loan, or public utility bill payments, verifying payment dates and amounts due". By contrast, as explained above, the workstyle data which we use in our measures are lower-dimensional and score all occupations on the same set of characteristics.

²³The weight $q_{i,k}$ reflects the frequency of a term within a document relative to its frequency across all documents, thereby emphasizing words that are distinctive to that document. We compute TF-IDF separately for the patent and occupation text corpora.

²⁴We use the `glove-wiki-gigaword-300` model, trained on the Wikipedia and Gigaword corpora, which provides 300-dimensional word vectors that preserve semantic relationships among words.

where θ_j^r is the routine/nonroutine share of tasks for occupation j , \mathcal{P}_t is the set of breakthrough patents granted in year t , and $\pi_{p,n}$ is the probability weight linking patent p to industry n from [Goldschlag, Lybbert, and Zolas \(2016\)](#).²⁵

Predict future employment of occupation j in industry n We then follow [Kogan et al. \(2024\)](#) and use occupation exposure to technologies $\xi_{j,n,t}^R$ and $\xi_{j,n,t}^{NR}$ to predict future employment. Specifically, we run a regression of log employment of occupation j in industry n and year $t + 5$ ($\log l_{j,n,t+5}$) on the occupation's exposure to year t technologies through routine tasks ($\xi_{j,n,t}^R$), the occupation's exposure to year t technologies through nonroutine tasks ($\xi_{j,n,t}^{NR}$), and current employment of occupation j in industry n ($\log l_{j,n,t}$). We obtain predicted log employment of occupation j in industry n and year $t + 5$ ($\widehat{\log l_{j,n,t+5}}$):

$$\widehat{\log l_{j,n,t+5}} = \underset{(0.060)}{0.308^{***}} - \underset{(0.009)}{0.052^{***}} \cdot \xi_{j,n,t}^R + \underset{(0.004)}{0.031^{***}} \cdot \xi_{j,n,t}^{NR} + \underset{(0.006)}{0.953^{***}} \cdot \log l_{j,n,t}. \quad (37)$$

This regression shows that the employment of an occupation declines by more when new technologies are more similar to its routine tasks, indicating substitution and replacement, while the employment rises by more when new technologies are more similar to its nonroutine tasks, indicating complementarity and enhancement. These effects align with intuition and appear robust in the data.

Predicted workstyle change in industry n from year t to $t + 5$ Finally, defining $\widehat{l_{n,t+5}} \equiv \sum_j \widehat{l_{j,n,t+5}}$, we obtain the predicted workstyle in year $t + 5$ using $\log \widehat{l_{j,n,t+5}}$ (e as expected or predicted):

$$\mathbf{WS}_{n,t+5}^e = \sum_{j=1}^J \frac{\widehat{l_{j,n,t+5}}}{\widehat{l_{n,t+5}}} \mathbf{ws}_j. \quad (38)$$

Accordingly, the predicted workstyle change in industry n from year t to $t + 5$:

$$|d\overline{ws}|_{n,t,5}^e \equiv \|\mathbf{WS}_{n,t+5}^e - \mathbf{WS}_{n,t}\|. \quad (39)$$

This predicted industry-level workstyle change based on technologies, $|d\overline{ws}|_{n,t,5}^e$, will be the key independent variable in our analysis in Section [V](#).

²⁵The [Goldschlag, Lybbert, and Zolas \(2016\)](#) bridge maps each Cooperative Patent Classification (CPC) code to an industry with a probability weight. We follow [Kelly et al. \(2021\)](#) to make the probability weights associated with a patent sum to one.

V Empirical Results

In this section, we analyze how workstyle changes in an industry help us understand the growth of young vs old firms. The key prediction is that such changes are more difficult for old firms to accommodate, and they grow more slowly than young firms when new technologies induce greater shifts in workstyle. We test this hypothesis in several datasets. In Section V.A, we measure the strength of young firms using venture capital investment. In Section V.B, we measure the performance of young firms vs old firms using Compustat data, which allow for a variety of metrics including valuation, sales growth, and employment growth. In Section V.C, we measure the growth of young firms vs old firms using employment growth in Census Business Dynamic Statistic (BDS) data. The BDS dataset offers broad coverage of the economy, although many young firms in that case may be subsistence entrepreneurship rather than innovative entrepreneurship. In addition, firms' exact age is unknown for those formed before the start of the BDS dataset in 1976. All of our tests use three-digit nonfarm nonfinancial NAICS industries (i.e., excluding NAICS codes starting with 11, 52, 53, and 55).²⁶ Since we rely on BLS employment composition data at the three-digit NAICS code level and O*NET workstyle data, our sample for the empirical analyses starts in 2003 when these data became consistently available; our sample ends in 2016 because several of the patent controls are available through 2016.

V.A Venture Capital Investment

We start with measuring the strength of young firms through venture capital (VC) investment, which captures forward-looking valuation of startups that can reflect their growth potential. This test assumes that VC investors understand conditions that favor the growth of young firms (e.g., through examining their competitive advantages).

We proceed as follows. First, we obtain venture capital investment from Refinitiv and calculate total VC investment in year t and industry n . We obtain a list of deals from Refinitiv's Private Equity database (Refinitiv, 2025a), and select Fund Investors Type to be Venture Capital; we then use a fuzzy matching procedure to standardize industry names in Refinitiv data (which are 2007 NAICS codes in text) and map them to numerical 2007 NAICS codes. We sum over VC investment value by three-digit NAICS and year to obtain the industry-year level VC investment.

In Table 1, we regress log VC investment in industry n following year t ($\log VC_{n,t+1}$) on $|d\overline{w}_s|_{n,t,5}^e$,

²⁶Since the Goldschlag, Lybbert, and Zolas (2016) crosswalk of patent technology classes to NAICS codes does not include 483 (water transportation) and 813 (religious, grantmaking, civic, professional, and similar organizations), our analyses also exclude these industries.

the predicted workstyle change over the next five years due to year t technologies. We use the dependent variable in year $t + 1$ to ensure that the year t predicted workstyle change is within the information set of investors. We also include year fixed effects to remove fluctuations in VC investment due to macroeconomic conditions.

$$\log VC_{n,t+1} = \alpha_t + \beta |d\overline{ws}|_{n,t,5}^e + \gamma z_{n,t} + \epsilon_{n,t}. \quad (40)$$

Since $|d\overline{ws}|_{n,t,5}^e$ is a generated regressor based on Equations (37) and (39), we bootstrap standard errors in all of our regression analyses. In each bootstrap draw, we resample industries with replacement, and use the corresponding bootstrap weights to reestimate the technology-implied employment change by industry-occupation-year in Equation (37) and recalculate the predicted workstyle change in Equation (39), which is then used to rerun the final regression of interest in the bootstrap sample.

In Table 1 column (1), we see that VC investment value is significantly higher when an industry is hit by new technologies that are predicted to change its workstyle. A one standard deviation increase in $|d\overline{ws}|_{n,t,5}^e$ is associated with higher VC investment by over 0.4 log points, which is meaningful. In the industry-year level controls $z_{n,t}$, we always include the log total market capitalization of Compustat firms in industry n and year $t + 1$ to capture other factors that may affect the prospects of firms in the industry.

In columns (2) to (5), we control for the volume of new technologies in the industry using several patent-based measures from prior work, and show that the sheer quantity of new technologies does not display a significant relationship with the strength of young firms; these controls also have a limited impact on the regression coefficients on $|d\overline{ws}|_{n,t,5}^e$. In column (2), we control for the log number of total patents in industry n and year t . Column (3) controls for the log number of breakthrough patents in industry n and year t using data from Kelly et al. (2021). The breakthrough patents aim to capture transformative new technologies, which are distinct from patents that came before them and followed by many similar patents afterwards. This measure is available for patents granted through 2016; we use the measure based on backward and forward significance over five years (instead of the version based on backward and forward significance over ten years, which is only available for patents granted through 2011), and a patent counts as a breakthrough one if its score is in the top 10%. Column (4) controls for the log number of creative patents using data from Kalyani (2025) available for patents filed through 2016. The creativity scores for patents are based on the share of their technical bigrams that did not appear in any patents filed over the past five years, and creative patents are those with the top 10% creativity scores. Column (5) controls for the log number of total rapidly evolving technology (RETech) patents in industry n and year t , using data from Bowen, Frésard, and Hoberg (2023). Patents using words that are contemporaneously surging across the patent corpus receive a higher score, and we count RETech patents as those with scores in the top 10%.

Table 1 – Venture Capital Investment

	Forward 1-Year Log(VC)				
	(1)	(2)	(3)	(4)	(5)
Predicted Workstyle Change in Industry	0.475*** (0.155)	0.463*** (0.159)	0.404*** (0.153)	0.426*** (0.155)	0.415*** (0.156)
Log Market Value	0.736*** (0.153)	0.732*** (0.165)	0.645*** (0.176)	0.669*** (0.175)	0.633*** (0.174)
Log Patents in Industry		-0.008 (0.093)	0.146 (0.116)	0.112 (0.129)	0.172 (0.126)
Patent Type	-	All	Breakthrough	Creative	RETech
Year FE	X	X	X	X	X
Observations	748	748	748	748	748
Within R^2	0.278	0.278	0.296	0.289	0.296

Notes: This table presents regressions at the industry-year level in Equation (40). The outcome variable is log VC investment in industry n following year t ($\log VC_{n,t+1}$). The independent variables include the predicted workstyle change in industry n over the next 5 years due to year t technologies ($|\overline{dws}|_{n,t,5}^e$), constructed in Equation (39). We standardize $|\overline{dws}|_{n,t,5}^e$ to unit variance and zero median. We control for log total market capitalization of Compustat firms in industry n ($\log MV_{n,t+1}$). Column (2) controls for the log number of total patents in industry n and year t . Column (3) controls for the log number of breakthrough patents in industry n and year t using data from Kelly et al. (2021). Column (4) controls for the log number of creative patents in industry n and year t using data from Kalyani (2025). Column (5) controls for the log number of rapidly evolving patents (RETech) in industry n and year t using data from Bowen, Frésard, and Hoberg (2023). We include year fixed effects, and bootstrap standard errors. In each bootstrap draw, we resample industries with replacement, and use the corresponding bootstrap weights to reestimate the technology-implied employment change by industry-occupation-year in Equation (37) and recalculate the predicted workstyle change in Equation (39), which is then used to rerun the final regression of interest in the bootstrap sample. The coefficient is the bootstrap median. Asterisks denote significance levels (***=1%, **=5%, *=10%). Number of observations and within R^2 (without fixed effects) are based on the OLS regression. Sample years are 2003 to 2016.

V.B Growth of Young vs Old Firms in Compustat Data

We then examine the strength of young vs old firms using Compustat data (S&P Global Market Intelligence, 2025). We combine information from several sources to measure the age of Compustat firms, which is not directly available in Compustat. First, we obtain firms' IPO date from Compustat and incorporation date from Refinitiv (Refinitiv, 2025b). The limitation of these data is that firms do not always become public soon after their founding, and they may reincorporate later; these data are also available only for a subset of Compustat firms (Compustat only began to systematically record IPO dates for IPOs after 2002). Second, we use LLM to query the founding year of each Compustat name (based on their real time historical name), as explained in more detail in Appendix Section IA5. This information is reliable for well known companies, but can be less accurate for companies that are not well recognized or those that no longer exist. Taken together, we calculate firm age as the number of years since the earliest of these dates (we set age to missing if the earliest of these dates is after the earliest year of the company's record in

Compustat, which indicates that our founding year information may not be accurate).

In Table 2, we perform regressions using Tobin’s Q (calculated as the market value of equity plus book value of debt, divided by book assets) as the outcome variable. It captures forward-looking valuation of young vs old firms regarding their growth potential, analogous to the VC investment analyses in Section V.A but now within the set of public companies (rather than startups relative to public companies). This test also assumes that investors understand conditions under which young firms are especially powerful (e.g., through examining their competitive advantages). Specifically, the outcome variable is Tobin’s Q of firm i in industry n and year $t + 1$ (to make sure year t technologies are within information sets similar to the VC regressions in Equation (40)). The independent variables include the predicted workstyle change in industry n over the next five years due to year- t technologies ($|d\overline{ws}|_{n,t,5}^e$), and its interaction with log firm age. The controls include the patent variables in columns (2) to (5) of Table 1, as well as their interactions with log firm age. To differentiate the effects of age from firm size, we also control for firm size and its interactions with the predicted workstyle change (as well as its interactions with the patent controls).²⁷ We include year fixed effects, and bootstrap standard errors as explained earlier.

$$Q_{i,t+1} = \alpha_t + \zeta \left(|d\overline{ws}|_{n,t,5}^e \times \text{age}_{i,t} \right) + \kappa |d\overline{ws}|_{n,t,5}^e + \rho \text{age}_{i,t} + \gamma z_{n,t} + \epsilon_{i,t}. \quad (41)$$

In Tables 3 and 4, we perform similar regressions using realized sales growth and employment growth in the five years after t as the outcome variables. The growth rates $\Delta Y_{i,t \rightarrow t+5}$ are calculated following Davis, Haltiwanger, and Schuh (1992), to be consistent with the specification we need to use in Census Business Dynamics Statistics dataset later in Section V.C. It is defined as $\Delta Y_{i,t \rightarrow t+5} = (Y_{i,t+5} - Y_{i,t}) / (0.5 \times Y_{i,t+5} + 0.5 \times Y_{i,t})$, which avoids missing values if the starting level is zero or extreme values when the starting level is small. The independent variables include the predicted workstyle change in industry n over the next five years due to year- t technologies ($|d\overline{ws}|_{n,t,5}^e$), and its interaction with log firm age, patent controls interacted with log firm age, as well as firm size and the same set of interactions. We include year fixed effects and bootstrap standard errors.

$$\Delta Y_{i,t \rightarrow t+5} = \alpha_t + \zeta \left(|d\overline{ws}|_{n,t,5}^e \times \text{age}_{i,t} \right) + \kappa |d\overline{ws}|_{n,t,5}^e + \rho \text{age}_{i,t} + \gamma z_{n,t} + \epsilon_{i,t}. \quad (42)$$

Tables 2, 3, and 4 show that when $|d\overline{ws}|_{n,t,5}^e$ is high, older firms have significantly lower valuation, as well as lower sales and employment growth. Meanwhile, the number of patents $z_{n,t}$ (all, breakthrough,

²⁷Appendix Section IA6.3.4 shows that adding industry-level and firm-level controls does not change the economic interpretation of the interaction coefficient between predicted workstyle change and firm age described in Result 2. In the presence of additional controls, the estimated coefficient is the product of an industry-level constant, which becomes conditional on other industry-level controls, and the covariance of age and organizational frictions conditional on other firm-level observables.

creative, or rapidly evolving) in the industry does not have any significant interaction with firm age. The coefficients on $|\overline{dws}|_{n,t,5}^e$ and its interaction with firm age are also not affected by the controls on the quantity of patents.

A natural question is whether the results on firm age reflect a correlation between age and size (Akçigit and Kerr, 2018): perhaps young firms are more nimble in light of technology-induced style changes because they are small. In Tables 2, 3, and 4, we do not observe that size is negatively associated with how firm outcomes respond to workstyle change. The interaction between predicted workstyle change $|\overline{dws}|_{n,t,5}^e$ and firm size tends to be positive. It could be possible that, all else equal, larger firms can have more financial resources that help them with adjustment and adaptation.

Figures IA1 and IA2 visualize the implied coefficient on $|\overline{dws}|_{n,t,5}^e$ for each level of firm age, using the specification in column (1) of Tables 2, 3, and 4. We transform the coefficient on log firm age to that on firm age for ease of illustration.

V.C Growth of Young vs Old Firms in BDS Data

Finally, we turn to Census Business Dynamics Statistics (BDS), which has broader coverage of firms. Here the primary measure of size is employment (United States Census Bureau, 2023). In particular, the public use dataset provides total employment of firms by age group: 0, 1-5, 6-10, 11-15, 16-20, 21-25, and the remaining age groups cannot be consistently defined over our sample period because firms' precise age is unknown if they are born before 1976. Therefore we restrict to firms with age groups between 1-5 and 16-20 (which become age groups 6-10 to 21-25 five years later). We perform regressions at the industry-year-age group level:

$$\Delta \text{emp}_{i,t \rightarrow t+5} = \alpha_t + \zeta \left(|\overline{dws}|_{n,t,5}^e \times \text{age}_{i,t} \right) + \kappa |\overline{dws}|_{n,t,5}^e + \rho \text{age}_{i,t} + \gamma z_{n,t} + \epsilon_{i,t}. \quad (43)$$

The outcome variable is the growth rate of employment among firm age group i in industry n between year t and $t + 5$ (e.g., age group 1-5 becomes 6-10 after five years), calculated following Davis, Haltiwanger, and Schuh (1992). The independent variables include the predicted workstyle change in industry n over the next five years due to year- t technologies ($|\overline{dws}|_{n,t,5}^e$), and its interaction with age group dummies. The controls include the patent variables in columns (2) to (5) of Table 1, as well as their interactions with age group dummies. We include year fixed effects, and bootstrap standard errors. We see that when $|\overline{dws}|_{n,t,5}^e$ is high, employment growth is faster for the young firm groups, and less so for the older firm groups.

Table 2 – Tobin’s Q among Compustat Firms

	Tobin’s Q				
	(1)	(2)	(3)	(4)	(5)
Predicted Workstyle Change in Industry \times Log Age	-0.159*** (0.045)	-0.158*** (0.039)	-0.143*** (0.034)	-0.149*** (0.036)	-0.155*** (0.039)
Predicted Workstyle Change in Industry \times Log Size	0.043* (0.020)	0.046** (0.018)	0.049** (0.019)	0.047** (0.018)	0.049** (0.020)
Predicted Workstyle Change in Industry	0.304* (0.173)	0.229 (0.154)	0.150 (0.147)	0.178 (0.145)	0.189 (0.149)
Log Patents in Industry \times Log Age		-0.004 (0.027)	-0.016 (0.030)	-0.017 (0.033)	-0.011 (0.032)
Log Patents in Industry \times Log Size		-0.006 (0.013)	-0.008 (0.018)	-0.009 (0.019)	-0.015 (0.019)
Log Patents in Industry		0.087 (0.147)	0.161 (0.145)	0.189 (0.167)	0.222 (0.155)
Log Age	-0.258*** (0.072)	-0.233 (0.169)	-0.189* (0.117)	-0.187* (0.120)	-0.203** (0.101)
Log Size	-0.193*** (0.038)	-0.151* (0.074)	-0.153** (0.059)	-0.148** (0.059)	-0.132** (0.051)
Patent Type	-	All	Breakthrough	Creative	RETech
Year FE	X	X	X	X	X
Observations	27,978	27,978	27,978	27,978	27,978
Within R^2	0.077	0.080	0.079	0.080	0.082

Notes: This table presents regressions at the firm-year level in Equation (41). The independent variables include the predicted workstyle change in industry n over the next 5 years due to year t technologies ($|d\bar{w}s|_{n,t,5}^e$, constructed in Equation (39), and its interaction with firm age and firm size (log book assets in year t). We standardize $|d\bar{w}s|_{n,t,5}^e$ to unit variance and zero median. Column (2) controls for the log number of total patents in industry n and year t , as well as its interactions with firm age and firm size. Column (3) controls for the log number of breakthrough patents in industry n and year t using data from Kelly et al. (2021), as well as its interactions with firm age and firm size. Column (4) controls for the log number of creative patents in industry n and year t using data from Kalyani (2025), as well as its interactions with firm age and firm size. Column (5) controls for the log number of rapidly evolving patents (RETech) in industry n and year t using data from Bowen, Frésard, and Hoberg (2023), as well as its interactions with firm age and firm size. We include year fixed effects, and bootstrap standard errors. In each bootstrap draw, we resample industries with replacement, and use the corresponding bootstrap weights to reestimate the technology-implied employment change by industry-occupation-year in Equation (37) and recalculate the predicted workstyle change in Equation (39), which is then used to rerun the final regression of interest in the bootstrap sample. The coefficient is the bootstrap median. Asterisks denote significance levels (***=1%, **=5%, *=10%). Number of observations and within R^2 (without fixed effects) are based on the OLS regression. Sample years are 2003 to 2016.

V.D Additional Tests

We perform several additional analyses in this section to further address alternative mechanisms.

Uncertainty Recent work by Goldman, Slutzky, and Zeume (2025) using global data finds that uncertainty appears to hurt old firms more than young firms: older firms contract significantly more than younger firms during periods of high uncertainty. In Tables IA7, IA8, IA9, we check that our results on

Table 3 – Sales Growth among Compustat Firms

	Sales Growth (5-Year DHS Rate)				
	(1)	(2)	(3)	(4)	(5)
Predicted Workstyle Change in Industry \times Log Age	-0.025** (0.009)	-0.026** (0.009)	-0.024** (0.009)	-0.025** (0.009)	-0.024** (0.010)
Predicted Workstyle Change in Industry \times Log Size	0.010* (0.005)	0.011** (0.006)	0.009* (0.005)	0.010* (0.006)	0.010* (0.006)
Predicted Workstyle Change in Industry	0.017 (0.031)	0.016 (0.038)	0.023 (0.036)	0.018 (0.038)	0.015 (0.036)
Log Patents in Industry \times Log Age		0.003 (0.007)	0.001 (0.009)	0.002 (0.009)	0.001 (0.009)
Log Patents in Industry \times Log Size		-0.001 (0.003)	0.002 (0.004)	-0.001 (0.005)	-0.001 (0.005)
Log Patents in Industry		-0.003 (0.024)	-0.010 (0.025)	-0.002 (0.029)	0.005 (0.026)
Log Age	-0.097*** (0.018)	-0.119*** (0.043)	-0.101*** (0.032)	-0.106*** (0.032)	-0.100*** (0.027)
Log Size	-0.025** (0.012)	-0.014 (0.015)	-0.027** (0.013)	-0.020 (0.012)	-0.020** (0.010)
Patent Type	-	All	Breakthrough	Creative	RETech
Year FE	X	X	X	X	X
Observations	27,978	27,978	27,978	27,978	27,978
Within R^2	0.059	0.059	0.059	0.059	0.059

Notes: This table presents regressions at the firm-year level in Equation (42), where the outcome variable $\Delta Y_{i,t \rightarrow t+5}$ is sales growth between year t and $t + 5$ calculated following [Davis, Haltiwanger, and Schuh \(1992\)](#) as $\Delta Y_{i,t \rightarrow t+5} = (Y_{i,t+5} - Y_{i,t}) / (0.5 \times Y_{i,t+5} + 0.5 \times Y_{i,t})$. The independent variables include the predicted workstyle change in industry n over the next 5 years due to year t technologies ($|d\bar{w}s|_{n,t,5}^e$), constructed in Equation (39), and its interaction with firm age and firm size (log sales in year t). We standardize $|d\bar{w}s|_{n,t,5}^e$ to unit variance and zero median. Column (2) controls for the log number of total patents in industry n and year t , as well as its interactions with firm age and firm size. Column (3) controls for the log number of breakthrough patents in industry n and year t using data from [Kelly et al. \(2021\)](#), as well as its interactions with firm age and firm size. Column (4) controls for the log number of creative patents in industry n and year t using data from [Kalyani \(2025\)](#), as well as its interactions with firm age and firm size. Column (5) controls for the log number of rapidly evolving patents (RETech) in industry n and year t using data from [Bowen, Frésard, and Hoberg \(2023\)](#), as well as its interactions with firm age and firm size. We include year fixed effects, and bootstrap standard errors. In each bootstrap draw, we resample industries with replacement, and use the corresponding bootstrap weights to reestimate the technology-implied employment change by industry-occupation-year in Equation (37) and recalculate the predicted workstyle change in Equation (39), which is then used to rerun the final regression of interest in the bootstrap sample. The coefficient is the bootstrap median. Asterisks denote significance levels (***=1%, **=5%, *=10%). Number of observations and within R^2 (without fixed effects) are based on the OLS regression. Sample years are 2003 to 2016.

predicted workstyle change are not driven by this variable being correlated with uncertainty. We use uncertainty measured from World Uncertainty Index like [Goldman, Slutzky, and Zeume \(2025\)](#), as well as economic policy uncertainty ([Baker, Bloom, and Davis, 2016](#)) and firm-level stock return volatility. The uncertainty controls do not affect our main results.

Table 4 – Employment Growth among Compustat Firms

	Employment Growth (5-Year DHS Rate)				
	(1)	(2)	(3)	(4)	(5)
Predicted Workstyle Change in Industry \times Log Age	-0.021** (0.008)	-0.020** (0.008)	-0.020** (0.008)	-0.020** (0.008)	-0.018** (0.008)
Predicted Workstyle Change in Industry \times Log Size	0.009* (0.004)	0.007** (0.003)	0.006** (0.003)	0.007** (0.003)	0.007** (0.003)
Predicted Workstyle Change in Industry	0.076** (0.027)	0.072** (0.030)	0.070** (0.031)	0.069* (0.031)	0.065* (0.031)
Log Patents in Industry \times Log Age		0.000 (0.006)	-0.001 (0.007)	-0.002 (0.007)	-0.004 (0.006)
Log Patents in Industry \times Log Size		0.001 (0.002)	0.002 (0.003)	0.002 (0.003)	0.002 (0.003)
Log Patents in Industry		-0.001 (0.021)	0.005 (0.024)	0.007 (0.026)	0.014 (0.023)
Log Age	-0.091*** (0.014)	-0.092** (0.035)	-0.087*** (0.025)	-0.086*** (0.026)	-0.081*** (0.022)
Log Size	-0.009 (0.006)	-0.013 (0.012)	-0.016* (0.010)	-0.014 (0.010)	-0.013* (0.008)
Patent Type	-	All	Breakthrough	Creative	RETech
Year FE	X	X	X	X	X
Observations	27,978	27,978	27,978	27,978	27,978
Within R^2	0.039	0.040	0.041	0.040	0.040

Notes: This table presents regressions at the firm-year level in Equation (42), where the outcome variable $\Delta Y_{i,t \rightarrow t+5}$ is employment growth between year t and $t + 5$ calculated following Davis, Haltiwanger, and Schuh (1992) as $\Delta Y_{i,t \rightarrow t+5} = (Y_{i,t+5} - Y_{i,t}) / (0.5 \times Y_{i,t+5} + 0.5 \times Y_{i,t})$. The independent variables include the predicted workstyle change in industry n over the next 5 years due to year t technologies ($|d\bar{w}s|_{n,t,5}^e$), constructed in Equation (39), and its interaction with firm age and firm size (log employment in year t). We standardize $|d\bar{w}s|_{n,t,5}^e$ to unit variance and zero median. Column (2) controls for the log number of total patents in industry n and year t , as well as its interactions with firm age and firm size. Column (3) controls for the log number of breakthrough patents in industry n and year t using data from Kelly et al. (2021), as well as its interactions with firm age and firm size. Column (4) controls for the log number of creative patents in industry n and year t using data from Kalyani (2025), as well as its interactions with firm age and firm size. Column (5) controls for the log number of rapidly evolving patents (RETech) in industry n and year t using data from Bowen, Frésard, and Hoberg (2023), as well as its interactions with firm age and firm size. We include year fixed effects, and bootstrap standard errors. In each bootstrap draw, we resample industries with replacement, and use the corresponding bootstrap weights to reestimate the technology-implied employment change by industry-occupation-year in Equation (37) and recalculate the predicted workstyle change in Equation (39), which is then used to rerun the final regression of interest in the bootstrap sample. The coefficient is the bootstrap median. Asterisks denote significance levels (***=1%, **=5%, *=10%). Number of observations and within R^2 (without fixed effects) are based on the OLS regression. Sample years are 2003 to 2016.

Cannibalization A classic hypothesis is that incumbents may fail to implement new technologies because doing so will cannibalize their existing businesses (Arrow, 1962). In particular, when existing businesses are highly profitable, the net gain from replacing it with new technologies would be smaller, which may weaken incumbents' incentives due to Arrow's replacement effect, even without any agency frictions. In other words, our core hypothesis is that incumbents face stronger organizational frictions in implementing new

Table 5 – Employment Growth among BDS Firms

	Employment Growth (5-Year DHS Rate)				
	(1)	(2)	(3)	(4)	(5)
Predicted Workstyle Change in Industry \times Age 6–10	0.005 (0.010)	0.004 (0.010)	0.004 (0.009)	0.004 (0.009)	0.004 (0.010)
Predicted Workstyle Change in Industry \times Age 11–15	-0.007 (0.018)	-0.007 (0.018)	-0.009 (0.018)	-0.008 (0.018)	-0.007 (0.018)
Predicted Workstyle Change in Industry \times Age 16–20	-0.026* (0.014)	-0.027* (0.014)	-0.029** (0.014)	-0.028* (0.014)	-0.028** (0.014)
Predicted Workstyle Change in Industry	0.021 (0.016)	0.022 (0.016)	0.023 (0.016)	0.022 (0.016)	0.021 (0.016)
Log Patents in Industry \times Age 6–10		0.004 (0.006)	0.005 (0.007)	0.005 (0.008)	0.007 (0.009)
Log Patents in Industry \times Age 11–15		0.007 (0.007)	0.012 (0.009)	0.009 (0.010)	0.004 (0.010)
Log Patents in Industry \times Age 16–20		0.014* (0.007)	0.017* (0.010)	0.015 (0.010)	0.015 (0.011)
Log Patents in Industry		-0.009 (0.009)	-0.010 (0.012)	-0.006 (0.012)	-0.001 (0.012)
Age 6–10	0.016 (0.017)	-0.004 (0.029)	0.004 (0.024)	0.005 (0.022)	0.004 (0.020)
Age 11–15	0.014 (0.021)	-0.015 (0.035)	-0.014 (0.032)	-0.004 (0.029)	0.008 (0.026)
Age 16–20	0.035 (0.023)	-0.028 (0.037)	-0.008 (0.032)	0.003 (0.030)	0.009 (0.027)
Patent Type	-	All	Breakthrough	Creative	RETech
Year FE	X	X	X	X	X
Observations	4,032	4,032	4,032	4,032	4,032
Within R^2	0.011	0.011	0.011	0.012	0.014

Notes: This table presents regressions at the industry-age group-year level in Equation (43), where the outcome variable $\Delta Y_{i,t \rightarrow t+5}$ is employment growth of age group i in industry n between year t and $t+5$, calculated following Davis, Haltiwanger, and Schuh (1992) as $\Delta Y_{i,t \rightarrow t+5} = (Y_{i,t+5} - Y_{i,t}) / (0.5 \times Y_{i,t+5} + 0.5 \times Y_{i,t})$. The independent variables include the predicted workstyle change in industry n over the next 5 years due to year t technologies ($|d\bar{w}s|_{n,t,5}^e$), constructed in Equation (39), and its interaction with age group dummies. We standardize $|d\bar{w}s|_{n,t,5}^e$ to unit variance and zero median. Column (2) controls for the log number of total patents in industry n and year t , as well as its interaction with firm age groups. Column (3) controls for the log number of breakthrough patents in industry n and year t using data from Kelly et al. (2021), as well as its interaction with firm age groups. Column (4) controls for the log number of creative patents in industry n and year t using data from Kalyani (2025), as well as its interaction with firm age groups. Column (5) controls for the log number of rapidly evolving patents (RETech) in industry n and year t using data from Bowen, Frésard, and Hoberg (2023), as well as its interaction with firm age groups. We include year fixed effects, and bootstrap standard errors. In each bootstrap draw, we resample industries with replacement, and use the corresponding bootstrap weights to reestimate the technology-implied employment change by industry-occupation-year in Equation (37) and recalculate the predicted workstyle change in Equation (39), which is then used to rerun the final regression of interest in the bootstrap sample. The coefficient is the bootstrap median. Asterisks denote significance levels (***=1%, **=5%, *=10%). Number of observations and within R^2 (without fixed effects) are based on the OLS regression. Sample years are 2003 to 2016.

technologies that require organizational change, which may exist independently of Arrow's replacement effect (and there is no obvious reason to expect any interaction effect between profitability and workstyle change). Nonetheless, it is useful to check that the role of firm age in our main results is not just driven by older firms having more current profits that may be cannibalized. In Appendix Tables [IA10](#), [IA11](#), [IA12](#), we perform additional analyses controlling for predicted workstyle change interacting with current profits, measured as net income over assets, net income over sales, and net income over stock market capitalization (effectively the inverse of P/E ratio). We verify that our main results are not affected by these controls. We also do not observe a clear interaction between profitability and predicted workstyle change.

VI Analysis Based on Wikipedia Titles

In Section [II](#), we used Wikipedia titles to present stylized facts about technologies implemented by young vs old firms. In this section, we return to Wikipedia titles as the measure of new technologies. In Section [VI.A](#), we repeat the main empirical analyses in Section [V](#) using Wikipedia titles instead of patents. In Section [VI.B](#), we further examine whether Wikipedia titles that imply greater workstyle change are more likely to be implemented by young firms, as measured in Section [II](#). The advantage of Wikipedia titles is that they capture technologies that are well-defined, so we can trace whether the implementation and commercialization are achieved by young vs old firms as in Section [II](#). This allows us to perform the additional test in Section [VI.B](#), and verify that new technologies associated with greater workstyle changes are more compatible with young firms (and less compatible with old firms). In comparison, patents are numerous and relatively micro, so it is rather difficult to trace out what types of firms implemented and commercialized the technologies in each patent.

VI.A Predicted Workstyle Change and Firm Outcomes using Wikipedia Titles

We first present the main analyses about the performance of young vs old firms using predicted workstyle change based on Wikipedia titles instead of patents. The data construction is similar to that in Section [IV.B](#), which we explain in Section [VI.A.1](#). We then present the results in Section [VI.A.2](#).

VI.A.1 Data Construction: Predicted Workstyle Change using Wikipedia Titles

We compute predicted workstyle change in each industry year using Wikipedia titles, analogous to the construction using patent data in Section [IV.B](#). In this case, we have 1,255 Wikipedia titles of technological

innovations in the U.S. between 2003 to 2016, to keep the sample coverage similar to the patent-based measures.

Calculate Wikipedia titles’ similarity with occupation j ’s routine and nonroutine tasks We start by extracting technical keywords from the Wikipedia title’s summary using LLM (Gemini 2.5 Pro), which capture the core mechanisms and characteristics of the title, as explained in detail in Appendix Section IA4.2. We represent each Wikipedia title w as the mean of its keywords’ word embeddings using the same *GloVe* model as in Section IV.B:

$$X_w = \frac{1}{K_w} \sum_{k=1}^{K_w} x_k, \quad (44)$$

where K_w is the number of keyword words for Wikipedia title w and x_k are *GloVe* word vectors. We use keywords for Wikipedia titles (instead of TF-IDF as in the case with patents in Section IV.B) because Wikipedia titles contain more background information (e.g., time and place of invention) aside from descriptions about the functions of the technology. The occupation task representations \mathcal{X}_j^r ($r \in \{R, NR\}$) are constructed the same way as in Section IV.B, using TF-IDF weighted *GloVe* embeddings of each occupation’s routine and nonroutine task descriptions (keeping only verbs and nouns and removing stop words before calculating the embeddings). We calculate cosine similarity $s_{w,j}^r$ between Wikipedia title w and the routine or nonroutine component of occupation j , analogous to (35), and apply the same adjustments of removing year fixed effects and imposing sparsity (at the 80th percentile) to obtain $\tilde{s}_{w,j}^r$.

Sum over Wikipedia titles in industry n and year t to get occupation j ’s exposure $\xi_{j,n,t}^R$ and $\xi_{j,n,t}^{NR}$ We assign each Wikipedia title to industries using semantic similarity.²⁸ Specifically, we represent each three-digit NAICS industry description as a TF-IDF weighted *GloVe* embedding, and compute cosine similarity between each Wikipedia title’s keyword embedding X_w and each NAICS industry embedding. We normalize the resulting similarities to probability weights $\pi_{w,n}$ that sum to one across industries for each technology (we impose sparsity by setting weights below the 80th percentile to zero before calculating the probability weights). The exposure of industry n and occupation j to technology at year t is then:

$$\xi_{j,n,t}^r = \theta_j^r \log \left(1 + \sum_{w \in \mathcal{W}_t} \pi_{w,n} \tilde{s}_{w,j}^r \right), \quad r \in \{R, NR\}, \quad (45)$$

²⁸We obtain three-digit NAICS industry descriptions from the 2017 NAICS classification (U.S. Census Bureau, 2017) and clean them using LLM (Gemini 2.5 Pro) to remove clauses describing what the industry does *not* include, retaining only descriptions of what the industry covers. We preprocess the cleaned NAICS descriptions using the same text pipeline as for the occupation tasks (lowercase, tokenize, POS-tag, keep nouns and verbs, remove stopwords, and lemmatize). Each NAICS industry is then represented as a TF-IDF-weighted *GloVe* embedding of its processed description. We compute cosine similarity between each Wikipedia title’s keyword embedding X_w and each NAICS industry embedding, impose sparsity by setting weights below the 80th percentile to zero, and normalize the remaining similarities to probability weights $\pi_{w,n}$ that sum to one across industries for each technology.

where θ_j^r is the routine/nonroutine share of tasks for occupation j , \mathcal{W}_t is the set of Wikipedia titles with invention year t , and $\pi_{w,n}$ is the probability weight linking Wikipedia title w to industry n .

Predict future employment of occupation j in industry n We follow the same approach as in Section IV.B and use occupation exposure to technologies $\xi_{j,n,t}^R$ and $\xi_{j,n,t}^{NR}$ from (45) to predict future employment $\widehat{\log l_{j,n,t+5}}$, by running the same regression as Equation (37) of log employment on exposure through routine tasks, exposure through nonroutine tasks, and current employment.

$$\widehat{\log l_{j,n,t+5}} = \underset{(0.061)}{0.314^{***}} - \underset{(0.137)}{0.749^{***}} \cdot \xi_{j,n,t}^R + \underset{(0.064)}{0.194^{**}} \cdot \xi_{j,n,t}^{NR} + \underset{(0.006)}{0.954^{***}} \cdot \log l_{j,n,t}. \quad (46)$$

Consistent with the results in (37) and with intuition, we observe that the employment of an occupation declines by more when new technologies are more similar to its routine tasks, indicating substitution and replacement, while the employment rises by more when new technologies are more similar to its nonroutine tasks, indicating complementarity and enhancement.

Predicted workstyle change in industry n from year t to $t + 5$ The predicted workstyle change $|\overline{dws}|_{n,t,5}^e$ is computed from the predicted employment $\widehat{\log l_{j,n,t+5}}$ following the same procedure as in Equations (38) and (39). As before, we standardize $|\overline{dws}|_{n,t,5}^e$ to unit variance and zero median.

VI.A.2 Empirical Results

We perform the analyses in Section V now using the predicted workstyle change in (46). Table IA2 repeats the analysis in Table 1; Tables IA3, IA4, and IA5 repeat the analyses in Tables 2, 3, and 4; Table IA6 repeats the analysis in Table 5. The results are very similar.

VI.B Predicted Workstyle Change and Technology Implementation

The analyses in Sections V and VI.A examine how young and old firms perform differently in industries experiencing greater predicted workstyle change. Wikipedia titles allow a further test: for each individual technology in a Wikipedia title, we can ask whether technologies that imply greater workstyle disruption are more likely to be implemented by young firms, as identified in Section II.

We construct predicted workstyle change associated with each individual Wikipedia technology w in its primary industry (the industry with the highest probability weight $\pi_{w,n}$ from the mapping in Section VI.A.1). We follow the same procedure as in Section VI.A.1, but compute the workstyle change attributable to technology w alone, rather than summing over all technologies in the industry year. Specifically, we

calculate occupation j 's exposure to technology w as $\theta_j^r \log(1 + \tilde{s}_{w,j}^r)$ for $r \in \{R, NR\}$ (where θ_j^r is the routine/nonroutine share of tasks r for occupation j and $\tilde{s}_{w,j}^r$ is the similarity between Wikipedia title w and the task), similar to Equation (45) except for a single technology. We then use the estimated coefficients from Equation (46) to derive predicted employment in five years $\hat{l}_{j,n,t+5}^w$ for each occupation associated with technology w , and in turn the associated workstyle $\mathbf{WS}_{n,t+5}^{e,w}$. The technology-level predicted workstyle change is then:

$$|d\overline{ws}|_{n,t,5}^{e,w} \equiv \|\mathbf{WS}_{n,t+5}^{e,w} - \mathbf{WS}_{n,t}\|. \quad (47)$$

where $\mathbf{WS}_{n,t}$ is the current workstyle in technology w 's primary industry n . We standardize $|d\overline{ws}|_{n,t,5}^{e,w}$ to unit variance and zero median.

We estimate the following technology-level regression:

$$\mathbf{1}[\text{Implemented by Young}_w] = \alpha_{t_w} + \beta |d\overline{ws}|_{n,t,5}^{e,w} + \epsilon_w, \quad (48)$$

where the independent variable is an indicator equal to one if the company most successful in the initial implementation and commercialization of the technology is a young firm (less than ten years old at the time), independently classified using LLM as described in Section II.A and Appendix Section IA2.3. The coefficient β captures whether technologies with greater predicted workstyle change are more likely to be implemented by young firms. We use U.S. technologies with invention years from 2003 to 2016 like in Section VI.A. We include year fixed effects based on invention year of technology w (in case Wikipedia coverage of technologies varies over time). We bootstrap standard errors as before. In each bootstrap draw, we resample industries with replacement, and use the corresponding bootstrap weights to reestimate the technology-level predicted employment change and recalculate the predicted workstyle change in Equation (47), which is then used to rerun the final regression of interest in the bootstrap sample.

Table 6 presents the results. We find that β is consistently positive and significant: technologies that predict larger workstyle changes are more likely to be implemented by young firms. Columns (1) and (2) include all technologies; Columns (3) and (4) restrict to technologies with definitive classifications (young or old), excluding those classified as ambiguous or never commercially implemented. Columns (2) and (4) additionally control for the length of the entire Wikipedia article (not just the summary), to account for the prominence or complexity of the technology. For example, one alternative hypothesis can be that more powerful technologies may lead to greater workstyle change, and they tend to be implemented by young firms. In the data, longer Wikipedia articles tend to be weakly associated with implementation by old firms instead (not statistically significant), and this control does not affect our key coefficient of interest on the predicted workstyle change. These technology-level findings provide direct evidence that

Table 6 – Technology Implementation by Young vs Old Firms

	1 [Implemented by Young]			
	All Technologies		Classified Technologies	
	(1)	(2)	(3)	(4)
Technology-level Predicted Workstyle Change	0.036*** (0.012)	0.036*** (0.012)	0.049*** (0.017)	0.048*** (0.017)
Log Wikipedia Article Length		-0.011 (0.017)		-0.050 (0.022)
Year FE	X	X	X	X
Observations	1,255	1,255	899	899
Within R^2	0.007	0.008	0.012	0.022

Notes: This table presents regressions at the technology level in Equation (48). The dependent variable is an indicator equal to one if the company most successful in the initial implementation and commercialization of the technology is a young firm (founded less than ten years before commercialization), classified using LLM as described in Section II.A and Appendix Section IA2.3. The key independent variable is the technology-level predicted workstyle change $|\overline{dws}|_{n,t,5}^{e,w}$ in Equation (47), standardized to unit variance and zero median. Columns (1) and (2) include all technologies; Columns (3) and (4) restrict to technologies with definitive firm type classifications, dropping those classified as ambiguous or never commercially implemented. All specifications include invention year fixed effects and use bootstrap standard errors. In each bootstrap draw, we resample industries with replacement, and use the corresponding bootstrap weights to reestimate the technology-level employment change (including reestimating coefficients from Equation (46)), and recalculate the predicted workstyle change in Equation (47), which is then used to rerun the final regression of interest in the bootstrap sample. The coefficient is the bootstrap median. Asterisks denote significance levels (***=1%, **=5%, *=10%). Number of observations and within R^2 (without fixed effects) are based on the OLS regression. Sample years are 2003 to 2016.

new technologies which require greater workstyle change are more likely to be implemented by young firms, consistent with the industry-level results.

VII Conclusion

Every year, over 200 NBER working papers are written on entrepreneurship ([National Bureau of Economic Research, 2025](#)). Given the popularity of this subject, it is essential to address the fundamental question: why don't old firms do new things, and when do we need new firms to implement new technologies? Our work highlights the importance of organizational frictions for understanding organizational change and disruption. In particular, existing firms are not always incapable of implementing new technologies due to the slowness of learning or fear of cannibalization. However, they are especially vulnerable when new technologies require changes in organizational processes and priorities, which naturally favor new organizations that start fresh. With advances in data and measurement—thanks to more ways to extract information from texts—future research can make further progress in capturing organizational frictions and investigating their role in innovation, productivity, and growth.

References

- Acemoglu, Daron, and David Autor.** 2011. "Skills, tasks and technologies: Implications for employment and earnings." In *Handbook of Labor Economics*. Vol. 4, 1043–1171. Elsevier.
- Acemoglu, Daron, Ufuk Akcigit, Harun Alp, Nicholas Bloom, and William Kerr.** 2018. "Innovation, reallocation, and growth." *American Economic Review*, 108(11): 3450–3491.
- Aghion, Philippe, and Jean Tirole.** 1994. "The management of innovation." *Quarterly Journal of Economics*, 109(4): 1185–1209.
- Aghion, Philippe, and Jean Tirole.** 1995. "Some implications of growth for organizational form and ownership structure." *European Economic Review*, 39(3-4): 440–455.
- Aghion, Philippe, and Jean Tirole.** 1997. "Formal and real authority in organizations." *Journal of Political Economy*, 105(1): 1–29.
- Aghion, Philippe, and Peter Howitt.** 1992. "A model of growth through creative destruction." *Econometrica*, 323–351.
- Akcigit, Ufuk, and William R Kerr.** 2018. "Growth through heterogeneous innovations." *Journal of Political Economy*, 126(4): 1374–1443.
- Anton, James J, and Dennis A Yao.** 1995. "Start-ups, spin-offs, and internal projects." *Journal of Law, Economics, and Organization*, 11(2): 362–378.
- Arrow, Kenneth J.** 1964. "Control in large organizations." *Management Science*, 10(3): 397–408.
- Arrow, Kenneth Joseph.** 1962. "Economic welfare and the allocation of resources for invention." In *Readings in Industrial Economics: Volume two: Private Enterprise and State Intervention*. 219–236. Springer.
- Asirvatham, Hemanth.** 2024. "The dynamo is not the computer." Working Paper.
- Atalay, Enghin, Ali Hortaçsu, and Chad Syverson.** 2014. "Vertical integration and input flows." *American Economic Review*, 104(4): 1120–1148.
- Autor, David H, Frank Levy, and Richard J Murnane.** 2003. "The skill content of recent technological change: An empirical exploration." *Quarterly Journal of Economics*, 118(4): 1279–1333.
- Baker, Scott R, Nicholas Bloom, and Steven J Davis.** 2016. "Measuring economic policy uncertainty." *Quarterly Journal of Economics*, 131(4): 1593–1636.
- Beaumont, Paul, Camille Hebert, and Victor Lyonnet.** 2025. "Build or buy? Human capital and corporate diversification." *Review of Financial Studies*, 38(5): 1333–1367.
- Bloom, Nicholas, and John Van Reenen.** 2007. "Measuring and explaining management practices across firms and countries." *Quarterly Journal of Economics*, 122(4): 1351–1408.
- Bowen, Donald E, Laurent Frésard, and Gerard Hoberg.** 2023. "Rapidly evolving technologies and startup exits." *Management Science*, 69(2): 940–967.
- Braguinsky, Serguey, Joonkyu Choi, Yuheng Ding, Karam Jo, and Seula Kim.** 2024. "Megafirms and recent trends in the US innovation: Empirical evidence from the US patent data." Working Paper.

- Bresnahan, Timothy F, Erik Brynjolfsson, and Lorin M Hitt.** 2002. "Information technology, workplace organization, and the demand for skilled labor: Firm-level evidence." *Quarterly Journal of Economics*, 117(1): 339–376.
- Caskurlu, Tolga, Gerard Hoberg, and Gordon M Phillips.** 2024. "New technology sectoral disruptions." Working Paper.
- Christensen, Clayton M.** 1997. *The Innovator's Dilemma: When New Technologies Cause Great Firms to Fail*. Harvard Business School Press.
- Coase, Ronald Harry.** 1937. "The nature of the firm." *Economica*, 4(16): 386–405.
- Cohan, William D.** 2022. *Power Failure: The Rise and Fall of an American Icon*. Penguin.
- Cohen, Lauren, Umit G Gurun, and Quoc H Nguyen.** 2022. "The ESG-innovation disconnect: Evidence from green patenting." Working Paper.
- Davis, River.** 2023. "The World's biggest carmaker made a huge bet on tech. Things went wrong fast." *The Wall Street Journal*. (Accessed November 9, 2025).
- Davis, Steven J., John Haltiwanger, and Scott Schuh.** 1992. "Gross job creation and destruction." *Journal of Labor Economics*, 10(3): 305–358.
- Decker, Ryan, John Haltiwanger, Ron Jarmin, and Javier Miranda.** 2014. "The role of entrepreneurship in US job creation and economic dynamism." *Journal of Economic Perspectives*, 28(3): 3–24.
- Dessein, Wouter, and Tano Santos.** 2006. "Adaptive organizations." *Journal of political Economy*, 114(5): 956–995.
- Dessein, Wouter, Luis Garicano, and Robert Gertner.** 2010. "Organizing for synergies." *American Economic Journal: Microeconomics*, 2(4): 77–114.
- Draghi, Mario.** 2024. "The future of European competitiveness - A competitiveness strategy for Europe.", https://commission.europa.eu/topics/eu-competitiveness/draghi-report_en.
- Ewens, Michael, and Matt Marx.** 2024. "Firm age and invention: An open-access dataset." Working Paper.
- Garcia-Macia, Daniel, Chang-Tai Hsieh, and Peter J Klenow.** 2019. "How destructive is innovation?" *Econometrica*, 87(5): 1507–1541.
- Garicano, Luis.** 2000. "Hierarchies and the organization of knowledge in production." *Journal of political economy*, 108(5): 874–904.
- Garicano, Luis, and Esteban Rossi-Hansberg.** 2004. "Inequality and the organization of knowledge." *American Economic Review*, 94(2): 197–202.
- George, Patrick.** 2022. "Volkswagen has given a name to its pain, and it is 'software.'" *The Verge*. (Accessed November 9, 2025).
- Gerstner, Louis V.** 2002. *Who Says Elephants Can't Dance? Inside IBM's Historic Turnaround*. HarperCollins Publishers.
- Goldman, Jim, Pablo Slutzky, and Stefan Zeume.** 2025. "A dark side to experience? Firm age and uncertainty." Working Paper.

- Goldschlag, Nathan, Travis J Lybbert, and Nikolas Jason Zolas.** 2016. "An 'algorithmic links with probabilities' crosswalk for uspc and cpc patent classifications with an application towards industrial technology composition." *US Census Bureau Center for Economic Studies paper no. CES-WP-16-15*.
- Gorton, Gary B, and Alexander K Zentefis.** 2024. "Corporate culture as a theory of the firm." *Economica*, 91(364): 1391–1423.
- Grossman, Sanford J, and Oliver D Hart.** 1986. "The costs and benefits of ownership: A theory of vertical and lateral integration." *Journal of Political Economy*, 94(4): 691–719.
- Hart, Oliver, and Bengt Holmstrom.** 2010. "A theory of firm scope." *Quarterly Journal of Economics*, 125(2): 483–513.
- Hart, Oliver, and John Moore.** 1990. "Property Rights and the Nature of the Firm." *Journal of political economy*, 98(6): 1119–1158.
- Hart, Oliver, and John Moore.** 2008. "Contracts as reference points." *Quarterly Journal of Economics*, 123(1): 1–48.
- Henderson, Rebecca M, and Kim B Clark.** 1990. "Architectural innovation: The reconfiguration of existing product technologies and the failure of established firms." *Administrative science quarterly*, 9–30.
- Holmstrom, Bengt.** 1989. "Agency costs and innovation." *Journal of Economic Behavior & Organization*, 12(3): 305–327.
- Holmstrom, Bengt, and Paul Milgrom.** 1991. "Multitask principal–agent analyses: Incentive contracts, asset ownership, and job design." *The Journal of Law, Economics, and Organization*, 7(special_issue): 24–52.
- Kalyani, Aakash.** 2025. "The creativity decline: Evidence from US patents." Working Paper.
- Kalyani, Aakash, Nicholas Bloom, Marcela Carvalho, Tarek Hassan, Josh Lerner, and Ahmed Tahoun.** 2025. "The diffusion of new technologies." *Quarterly Journal of Economics*, 140(2): 1299–1365.
- Kelly, Bryan, Dimitris Papanikolaou, Amit Seru, and Matt Taddy.** 2021. "Measuring technological innovation over the long run." *American Economic Review: Insights*, 3(3): 303–320.
- Kogan, Leonid, Dimitris Papanikolaou, Lawrence DW Schmidt, and Bryan Seegmiller.** 2024. "Technology and labor displacement: Evidence from linking patents with worker-level data." Working Paper.
- Lee, Kyeong Hun, David C Mauer, and Emma Qianying Xu.** 2018. "Human capital relatedness and mergers and acquisitions." *Journal of Financial Economics*, 129(1): 111–135.
- Leonard-Barton, Dorothy.** 1992. "Core capabilities and core rigidities: A paradox in managing new product development." *Strategic Management Journal*, 13(S1): 111–125.
- Levitt, Steven D, John A List, and Chad Syverson.** 2013. "Toward an understanding of learning by doing: Evidence from an automobile assembly plant." *Journal of Political Economy*, 121(4): 643–681.
- Loderer, Claudio, René Stulz, and Urs Waelchli.** 2017. "Firm rigidities and the decline in growth opportunities." *Management Science*, 63(9): 3000–3020.
- March, James G.** 1991. "Exploration and exploitation in organizational learning." *Organization Science*, 2(1): 71–87.
- Meyer, Margaret, Paul Milgrom, and John Roberts.** 1992. "Organizational prospects, influence costs, and ownership changes." *Journal of Economics & Management Strategy*, 1(1): 9–35.

- Milgrom, Paul R, and John Roberts.** 1992. *Economics, Organization and Management*. Vol. 7, Prentice-hall Englewood Cliffs, NJ.
- Minni, Virginia M. L.** 2026. "Making the invisible hand visible: Managers and the allocation of workers to jobs." *Quarterly Journal of Economics*. Forthcoming.
- National Bureau of Economic Research.** 2025. "NBER working papers and chapters metadata." <https://www.nber.org/research/data/nber-working-papers-and-chapters-metadata> Updated weekly; available in CSV, TSV, and Stata (DTA) formats.
- Ndiaye, Pap.** 2007. *Nylon and Bombs: DuPont and the March of Modern America*. JHU Press.
- Nelson, Richard R, and Sidney G Winter.** 1985. *An evolutionary theory of economic change*. Harvard University Press.
- O'Reilly, Charles A, and Michael L Tushman.** 2021. *Lead and Disrupt: How to Solve the Innovator's Dilemma*. Stanford University Press.
- Penrose, Edith T.** 1959. *The Theory of the Growth of the Firm*. Oxford:Basil Blackwell.
- Qian, Yingyi, Gerard Roland, and Chenggang Xu.** 2006. "Coordination and experimentation in M-form and U-form organizations." *Journal of Political Economy*, 114(2): 366–402.
- Rajan, Raghuram G.** 2012. "Presidential address: The corporation in finance." *Journal of Finance*, 67(4): 1173–1217.
- Rajan, Raghuram G, and Luigi Zingales.** 2001. "The firm as a dedicated hierarchy: A theory of the origins and growth of firms." *Quarterly Journal of Economics*, 116(3): 805–851.
- Rajan, Raghuram, Henri Servaes, and Luigi Zingales.** 2000. "The cost of diversity: The diversification discount and inefficient investment." *Journal of Finance*, 55(1): 35–80.
- Refinitiv.** 2025a. "Refinitiv / LSEG workspace private equity screener." <https://workspace.refinitiv.com/web/Apps/ScreeningApp/?universe=PVTInvestment> Accessed through Refinitiv LSEG (Refinitiv Workspace) platform.
- Refinitiv.** 2025b. "Refinitiv Eikon / Workspace Database." <https://www.refinitiv.com/en/products/eikon-trading-software> Accessed through Refinitiv Eikon (Refinitiv Workspace) platform.
- Revelio Labs.** 2025. "Individual reviews (sentiment) dataset." <https://wrds-www.wharton.upenn.edu/pages/get-data/revelio-labs/revelio-sentiment/individual-reviews/> Accessed November 6, 2025.
- Scharfstein, David S, and Jeremy C Stein.** 2000. "The dark side of internal capital markets: Divisional rent-seeking and inefficient investment." *Journal of Finance*, 55(6): 2537–2564.
- Schumpeter, Joseph Alois.** 1942. *Capitalism, Socialism and Democracy*. Allen & Unwin London.
- Shleifer, Andrei, and Robert W Vishny.** 1989. "Management entrenchment: The case of manager-specific investments." *Journal of Financial Economics*, 25(1): 123–139.
- S&P Global Market Intelligence.** 2025. "Compustat North America Fundamentals Annual Database." <https://wrds-www.wharton.upenn.edu/> Accessed via Wharton Research Data Services (WRDS).

- Stein, Jeremy C.** 1997. "Internal capital markets and the competition for corporate resources." *Journal of Finance*, 52(1): 111–133.
- Teece, David J, Gary Pisano, and Amy Shuen.** 1997. "Dynamic capabilities and strategic management." *Strategic Management Journal*, 18(7): 509–533.
- Tushman, Michael L, and Philip Anderson.** 1986. "Technological discontinuities and organizational environments." *Administrative Science Quarterly*, 31(3): 439–65.
- United States Census Bureau.** 2023. "Business Dynamics Statistics Datasets." <https://www.census.gov/data/datasets/time-series/econ/bds/bds-datasets.html> Accessed September 25, 2025".
- U.S. Bureau of Labor Statistics.** 2025. "Occupational Employment and Wage Statistics (OEWS) Tables." <https://www.bls.gov/oes/tables.htm> Accessed November 1, 2025.
- U.S. Census Bureau.** 2017. "2017 NAICS Descriptions." , https://www.census.gov/naics/2017NAICS/2017_NAICS_Descriptions.xlsx (Accessed February 2, 2026).
- U.S. Department of Labor, Employment and Training Administration.** 2005. "O*NET Database, Version 9.0 (December 2005)." https://www.onetcenter.org/dl_files/ncsc/onet90ac.zip Accessed November 8, 2025.
- U.S. Department of Labor, Employment and Training Administration.** 2025. "O*NET® Database Releases Archive." https://www.onetcenter.org/db_releases.html Accessed September 30, 2025.
- Williamson, Oliver E.** 1971. "The vertical integration of production: Market failure considerations." *American Economic Review*, 61(2): 112–123.

Internet Appendix

IA1 Additional Figures and Tables




Figure IA1. Tobin's Q among Compustat Firms: Implied Coefficient on $|\overline{dws}|_{n,t,5}^e$

Notes: This figure shows the implied coefficient on $|\overline{dws}|_{n,t,5}^e$ for different levels of firm age in Table 2, column (1).

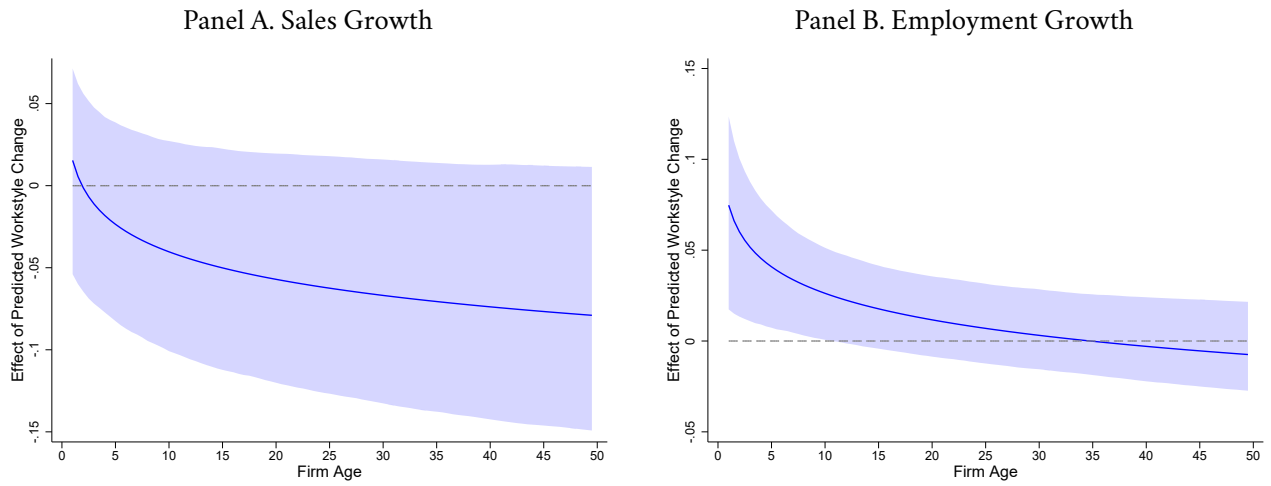


Figure IA2. Sales and Employment Growth among Compustat Firms: Implied Coefficient on $|dws|_{n,t,5}^e$

Notes: Panel A shows the implied coefficient on $|dws|_{n,t,5}^e$ for different levels of firm age in Table 3, column (1). Panel B shows the implied coefficient on $|dws|_{n,t,5}^e$ for different levels of firm age in Table 4, column (1).

Table IA1 – Summary Statistics

	P25	P50	P75	Mean	SD	N
Panel A: Venture Capital Sample						
Predicted Workstyle Change in Industry (Patent)	-0.15	-0.00	0.30	0.34	1.03	748
Predicted Workstyle Change in Industry (Wiki)	-0.13	-0.01	0.25	0.32	1.02	748
Log VC Investment in Industry	16.75	18.15	19.70	18.18	2.29	748
Log Compustat Market Cap in Industry	10.29	11.37	12.60	11.40	1.51	748
Log Total Patents in Industry	2.81	5.07	7.16	4.85	2.73	748
Log Breakthrough Patents in Industry	0.37	2.66	4.61	2.73	2.32	748
Log Creative Patents in Industry	0.53	2.35	4.03	2.44	2.01	748
Log RETech Patents in Industry	0.10	1.34	3.44	1.98	2.00	748
Panel B: Compustat Sample						
Predicted Workstyle Change in Industry (Patent)	-0.07	0.21	0.80	0.56	1.02	27,978
Predicted Workstyle Change in Industry (Wiki)	-0.11	0.08	0.59	0.44	0.97	27,978
Log Firm Age	2.89	3.47	4.19	3.51	0.87	27,978
Firm Tobin's Q	0.91	1.33	2.15	1.91	2.10	27,978
Sales Growth Next 5 Years	-0.03	0.25	0.58	0.27	0.56	27,978
Employment Growth Next 5 Years	-0.11	0.14	0.44	0.16	0.48	27,978
Log Firm Assets	4.52	6.13	7.65	6.08	2.26	27,978
Log Firm Sales	4.30	6.09	7.56	5.85	2.48	27,978
Log Firm Employment	-1.31	0.45	1.93	0.33	2.25	27,978
Log Total Patents in Industry	3.81	7.60	8.81	6.33	3.21	27,978
Log Breakthrough Patents in Industry	1.15	4.86	6.31	4.25	2.91	27,978
Log Creative Patents in Industry	1.22	4.63	5.76	3.84	2.54	27,978
Log RETech Patents in Industry	0.78	3.63	5.66	3.43	2.60	27,978
Panel C: BDS Sample						
Predicted Workstyle Change in Industry (Patent)	-0.14	-0.00	0.30	0.33	1.01	4,032
Predicted Workstyle Change in Industry (Wiki)	-0.13	-0.00	0.25	0.33	1.01	4,032
Employment Growth Next 5 Years	-0.24	-0.09	0.02	-0.10	0.26	4,032
Log Employment in Industry	12.79	13.43	14.21	13.46	1.26	4,032
Log Total Patents in Industry	1.99	4.24	6.58	4.40	2.68	4,032
Log Breakthrough Patents in Industry	0.31	2.01	4.31	2.44	2.23	4,032
Log Creative Patents in Industry	0.33	1.65	3.45	2.13	1.93	4,032
Log RETech Patents in Industry	0.07	1.00	3.11	1.71	1.91	4,032

Notes: This table reports summary statistics for the main variables used in the analysis. Panel A presents statistics for the venture capital sample at the industry-year level. Panel B presents statistics for the Compustat sample at the firm-year level. Panel C presents statistics for the BDS sample at the industry-age group-year level. Reported values include the 25th, 50th, and 75th percentiles (P25, P50, P75), mean, and standard deviation (SD).

Table IA2 – Venture Capital Investment: Wikipedia Title Results

	Forward 1-Year Log(VC)				
	(1)	(2)	(3)	(4)	(5)
Predicted Workstyle Change in Industry	0.444*** (0.168)	0.429*** (0.170)	0.399** (0.164)	0.410** (0.166)	0.407** (0.169)
Log Market Value	0.751*** (0.154)	0.738*** (0.167)	0.646*** (0.177)	0.671*** (0.176)	0.634*** (0.175)
Log Patents in Industry		0.008 (0.094)	0.163 (0.114)	0.132 (0.127)	0.189 (0.125)
Patent Type	-	All	Breakthrough	Creative	RETech
Year FE	X	X	X	X	X
Observations	748	748	748	748	748
Within R^2	0.273	0.274	0.297	0.288	0.296

Notes: This table presents regressions at the industry-year level in Equation (40). The outcome variable is log VC investment in industry n following year t ($\log VC_{n,t+1}$). The independent variables include the predicted workstyle change in industry n over the next 5 years due to year t Wikipedia technologies ($|\overline{dws}|_{n,t,5}^e$), constructed in Equation (39). We standardize $|\overline{dws}|_{n,t,5}^e$ to unit variance and zero median. We control for log total market capitalization of Compustat firms in industry n ($\log MV_{n,t+1}$). Column (2) controls for the log number of total patents in industry n and year t . Column (3) controls for the log number of breakthrough patents in industry n and year t using data from Kelly et al. (2021). Column (4) controls for the log number of creative patents in industry n and year t using data from Kalyani (2025). Column (5) controls for the log number of rapidly evolving patents (RETech) in industry n and year t using data from Bowen, Frésard, and Hoberg (2023). We include year fixed effects, and bootstrap standard errors. In each bootstrap draw, we resample industries with replacement, and use the corresponding bootstrap weights to reestimate the technology-implied employment change by industry-occupation-year in Equation (46) and recalculate the predicted workstyle change in Equation (39), which is then used to rerun the final regression of interest in the bootstrap sample. The coefficient is the bootstrap median. Asterisks denote significance levels (***=1%, **=5%, *=10%). Number of observations and within R^2 (without fixed effects) are based on the OLS regression. Sample years are 2003 to 2016.

Table IA3 – Tobin’s Q among Compustat Firms: Wikipedia Title Results

	Tobin’s Q				
	(1)	(2)	(3)	(4)	(5)
Predicted Workstyle Change in Industry \times Log Age	-0.148*** (0.045)	-0.145*** (0.040)	-0.130*** (0.035)	-0.137*** (0.036)	-0.143*** (0.039)
Predicted Workstyle Change in Industry \times Log Size	0.044** (0.018)	0.044** (0.018)	0.046** (0.018)	0.045** (0.018)	0.045** (0.020)
Predicted Workstyle Change in Industry	0.219 (0.186)	0.181 (0.164)	0.105 (0.160)	0.148 (0.152)	0.173 (0.155)
Log Patents in Industry \times Log Age		-0.011 (0.029)	-0.025 (0.031)	-0.027 (0.034)	-0.021 (0.033)
Log Patents in Industry \times Log Size		-0.003 (0.013)	-0.005 (0.017)	-0.005 (0.018)	-0.011 (0.019)
Log Patents in Industry		0.099 (0.149)	0.173 (0.145)	0.202 (0.167)	0.234 (0.157)
Log Age	-0.279*** (0.076)	-0.206 (0.182)	-0.175 (0.123)	-0.170 (0.125)	-0.193** (0.105)
Log Size	-0.190*** (0.038)	-0.159* (0.074)	-0.157** (0.058)	-0.153** (0.059)	-0.135** (0.051)
Patent Type	-	All	Breakthrough	Creative	RETech
Year FE	X	X	X	X	X
Observations	27,978	27,978	27,978	27,978	27,978
Within R^2	0.076	0.079	0.079	0.080	0.082

Notes: This table presents regressions at the firm-year level in Equation (41). The independent variables include the predicted workstyle change in industry n over the next 5 years due to year t Wikipedia technologies ($|d\bar{w}s|_{n,t,5}^e$), constructed in Equation (39), and its interaction with firm age and firm size (log book assets in year t). We standardize $|d\bar{w}s|_{n,t,5}^e$ to unit variance and zero median. Column (2) controls for the log number of total patents in industry n and year t , as well as its interactions with firm age and firm size. Column (3) controls for the log number of breakthrough patents in industry n and year t using data from Kelly et al. (2021), as well as its interactions with firm age and firm size. Column (4) controls for the log number of creative patents in industry n and year t using data from Kalyani (2025), as well as its interactions with firm age and firm size. Column (5) controls for the log number of rapidly evolving patents (RETech) in industry n and year t using data from Bowen, Frésard, and Hoberg (2023), as well as its interactions with firm age and firm size. We include year fixed effects, and bootstrap standard errors. In each bootstrap draw, we resample industries with replacement, and use the corresponding bootstrap weights to reestimate the technology-implied employment change by industry-occupation-year in Equation (46) and recalculate the predicted workstyle change in Equation (39), which is then used to rerun the final regression of interest in the bootstrap sample. The coefficient is the bootstrap median. Asterisks denote significance levels (***=1%, **=5%, *=10%). Number of observations and within R^2 (without fixed effects) are based on the OLS regression. Sample years are 2003 to 2016.

Table IA4 – Sales Growth among Compustat Firms: Wikipedia Title Results

	Sales Growth (5-Year DHS Rate)				
	(1)	(2)	(3)	(4)	(5)
Predicted Workstyle Change in Industry \times Log Age	-0.025** (0.009)	-0.024** (0.009)	-0.023** (0.009)	-0.024** (0.009)	-0.023** (0.010)
Predicted Workstyle Change in Industry \times Log Size	0.012** (0.006)	0.012** (0.006)	0.010* (0.006)	0.011* (0.006)	0.011* (0.006)
Predicted Workstyle Change in Industry	0.006 (0.037)	0.005 (0.040)	0.008 (0.040)	0.006 (0.041)	0.006 (0.039)
Log Patents in Industry \times Log Age		0.002 (0.007)	-0.001 (0.009)	0.001 (0.009)	-0.001 (0.009)
Log Patents in Industry \times Log Size		-0.001 (0.003)	0.002 (0.004)	0.000 (0.004)	-0.000 (0.004)
Log Patents in Industry		-0.002 (0.024)	-0.007 (0.025)	0.000 (0.028)	0.006 (0.026)
Log Age	-0.100*** (0.018)	-0.115** (0.044)	-0.099*** (0.032)	-0.103*** (0.032)	-0.098*** (0.027)
Log Size	-0.024** (0.012)	-0.016 (0.015)	-0.028** (0.013)	-0.021* (0.012)	-0.021** (0.010)
Patent Type	-	All	Breakthrough	Creative	RETech
Year FE	X	X	X	X	X
Observations	27,978	27,978	27,978	27,978	27,978
Within R^2	0.059	0.059	0.060	0.059	0.059

Notes: This table presents regressions at the firm-year level in Equation (42), where the outcome variable $\Delta Y_{i,t \rightarrow t+5}$ is sales growth between year t and $t + 5$ calculated following Davis, Haltiwanger, and Schuh (1992) as $\Delta Y_{i,t \rightarrow t+5} = (Y_{i,t+5} - Y_{i,t}) / (0.5 \times Y_{i,t+5} + 0.5 \times Y_{i,t})$. The independent variables include the predicted workstyle change in industry n over the next 5 years due to year t Wikipedia technologies ($|d\bar{w}s|_{n,t,5}^e$), constructed in Equation (39), and its interaction with firm age and firm size (log sales in year t). We standardize $|d\bar{w}s|_{n,t,5}^e$ to unit variance and zero median. Column (2) controls for the log number of total patents in industry n and year t , as well as its interactions with firm age and firm size. Column (3) controls for the log number of breakthrough patents in industry n and year t using data from Kelly et al. (2021), as well as its interactions with firm age and firm size. Column (4) controls for the log number of creative patents in industry n and year t using data from Kalyani (2025), as well as its interactions with firm age and firm size. Column (5) controls for the log number of rapidly evolving patents (RETech) in industry n and year t using data from Bowen, Frésard, and Hoberg (2023), as well as its interactions with firm age and firm size. We include year fixed effects, and bootstrap standard errors. In each bootstrap draw, we resample industries with replacement, and use the corresponding bootstrap weights to reestimate the technology-implied employment change by industry-occupation-year in Equation (46) and recalculate the predicted workstyle change in Equation (39), which is then used to rerun the final regression of interest in the bootstrap sample. The coefficient is the bootstrap median. Asterisks denote significance levels (***=1%, **=5%, *=10%). Number of observations and within R^2 (without fixed effects) are based on the OLS regression. Sample years are 2003 to 2016.

Table IA5 – Employment Growth among Compustat Firms: Wikipedia Title Results

	Employment Growth (5-Year DHS Rate)				
	(1)	(2)	(3)	(4)	(5)
Predicted Workstyle Change in Industry \times Log Age	-0.022** (0.008)	-0.020** (0.008)	-0.019** (0.008)	-0.019** (0.008)	-0.019** (0.008)
Predicted Workstyle Change in Industry \times Log Size	0.009** (0.003)	0.007** (0.003)	0.006** (0.003)	0.007** (0.003)	0.007** (0.003)
Predicted Workstyle Change in Industry	0.075** (0.027)	0.069* (0.030)	0.067* (0.031)	0.067* (0.031)	0.066* (0.031)
Log Patents in Industry \times Log Age		-0.001 (0.006)	-0.002 (0.006)	-0.003 (0.007)	-0.005 (0.006)
Log Patents in Industry \times Log Size		0.001 (0.002)	0.003 (0.003)	0.002 (0.003)	0.002 (0.003)
Log Patents in Industry		0.002 (0.021)	0.008 (0.023)	0.011 (0.026)	0.018 (0.023)
Log Age	-0.093*** (0.014)	-0.089** (0.035)	-0.085*** (0.025)	-0.084*** (0.026)	-0.080*** (0.022)
Log Size	-0.008 (0.006)	-0.014 (0.012)	-0.017* (0.010)	-0.014 (0.010)	-0.013* (0.008)
Patent Type	-	All	Breakthrough	Creative	RETech
Year FE	X	X	X	X	X
Observations	27,978	27,978	27,978	27,978	27,978
Within R^2	0.039	0.040	0.041	0.040	0.040

Notes: This table presents regressions at the firm-year level in Equation (42), where the outcome variable $\Delta Y_{i,t \rightarrow t+5}$ is employment growth between year t and $t + 5$ calculated following Davis, Haltiwanger, and Schuh (1992) as $\Delta Y_{i,t \rightarrow t+5} = (Y_{i,t+5} - Y_{i,t}) / (0.5 \times Y_{i,t+5} + 0.5 \times Y_{i,t})$. The independent variables include the predicted workstyle change in industry n over the next 5 years due to year t Wikipedia technologies ($|d\bar{w}s|_{n,t,5}^e$), constructed in Equation (39), and its interaction with firm age and firm size (log employment in year t). We standardize $|d\bar{w}s|_{n,t,5}^e$ to unit variance and zero median. Column (2) controls for the log number of total patents in industry n and year t , as well as its interactions with firm age and firm size. Column (3) controls for the log number of breakthrough patents in industry n and year t using data from Kelly et al. (2021), as well as its interactions with firm age and firm size. Column (4) controls for the log number of creative patents in industry n and year t using data from Kalyani (2025), as well as its interactions with firm age and firm size. Column (5) controls for the log number of rapidly evolving patents (RETech) in industry n and year t using data from Bowen, Frésard, and Hoberg (2023), as well as its interactions with firm age and firm size. We include year fixed effects, and bootstrap standard errors. In each bootstrap draw, we resample industries with replacement, and use the corresponding bootstrap weights to reestimate the technology-implied employment change by industry-occupation-year in Equation (46) and recalculate the predicted workstyle change in Equation (39), which is then used to rerun the final regression of interest in the bootstrap sample. The coefficient is the bootstrap median. Asterisks denote significance levels (***=1%, **=5%, *=10%). Number of observations and within R^2 (without fixed effects) are based on the OLS regression. Sample years are 2003 to 2016.

Table IA6 – Employment Growth among BDS Firms: Wikipedia Title Results

	Employment Growth (5-Year DHS Rate)				
	(1)	(2)	(3)	(4)	(5)
Predicted Workstyle Change in Industry \times Age 6–10	0.005 (0.009)	0.005 (0.009)	0.005 (0.009)	0.005 (0.009)	0.005 (0.009)
Predicted Workstyle Change in Industry \times Age 11–15	-0.009 (0.018)	-0.009 (0.018)	-0.009 (0.018)	-0.009 (0.018)	-0.009 (0.018)
Predicted Workstyle Change in Industry \times Age 16–20	-0.030** (0.014)	-0.028** (0.014)	-0.030** (0.014)	-0.030** (0.014)	-0.030** (0.014)
Predicted Workstyle Change in Industry	0.023 (0.016)	0.022 (0.016)	0.023 (0.016)	0.022 (0.016)	0.022 (0.016)
Log Patents in Industry \times Age 6–10		0.005 (0.006)	0.005 (0.007)	0.005 (0.008)	0.007 (0.009)
Log Patents in Industry \times Age 11–15		0.007 (0.007)	0.011 (0.009)	0.008 (0.010)	0.004 (0.010)
Log Patents in Industry \times Age 16–20		0.013* (0.007)	0.016* (0.010)	0.014 (0.010)	0.014 (0.011)
Log Patents in Industry		-0.008 (0.009)	-0.009 (0.012)	-0.005 (0.012)	-0.000 (0.012)
Age 6–10	0.017 (0.017)	-0.005 (0.029)	0.004 (0.025)	0.005 (0.023)	0.004 (0.020)
Age 11–15	0.015 (0.021)	-0.014 (0.035)	-0.014 (0.033)	-0.002 (0.029)	0.009 (0.026)
Age 16–20	0.036 (0.023)	-0.024 (0.038)	-0.005 (0.032)	0.006 (0.030)	0.012 (0.028)
Patent Type	-	All	Breakthrough	Creative	RETech
Year FE	X	X	X	X	X
Observations	4,032	4,032	4,032	4,032	4,032
Within R^2	0.011	0.011	0.011	0.011	0.014

Notes: This table presents regressions at the industry-age group-year level in Equation (43), where the outcome variable $\Delta Y_{i,t \rightarrow t+5}$ is employment growth of age group i in industry n between year t and $t+5$, calculated following Davis, Haltiwanger, and Schuh (1992) as $\Delta Y_{i,t \rightarrow t+5} = (Y_{i,t+5} - Y_{i,t}) / (0.5 \times Y_{i,t+5} + 0.5 \times Y_{i,t})$. The independent variables include the predicted workstyle change in industry n over the next 5 years due to year t Wikipedia technologies ($|d\bar{w}s|_{n,t,5}^e$), constructed in Equation (39), and its interaction with age group dummies. We standardize $|d\bar{w}s|_{n,t,5}^e$ to unit variance and zero median. Column (2) controls for the log number of total patents in industry n and year t , as well as its interaction with firm age groups. Column (3) controls for the log number of breakthrough patents in industry n and year t using data from Kelly et al. (2021), as well as its interaction with firm age groups. Column (4) controls for the log number of creative patents in industry n and year t using data from Kalyani (2025), as well as its interaction with firm age groups. Column (5) controls for the log number of rapidly evolving patents (RETech) in industry n and year t using data from Bowen, Frésard, and Hoberg (2023), as well as its interaction with firm age groups. We include year fixed effects, and bootstrap standard errors. In each bootstrap draw, we resample industries with replacement, and use the corresponding bootstrap weights to reestimate the technology-implied employment change by industry-occupation-year in Equation (46) and recalculate the predicted workstyle change in Equation (39), which is then used to rerun the final regression of interest in the bootstrap sample. The coefficient is the bootstrap median. Asterisks denote significance levels (***=1%, **=5%, *=10%). Number of observations and within R^2 (without fixed effects) are based on the OLS regression. Sample years are 2003 to 2016.

Table IA7 – Tobin’s Q among Compustat Firms, Controlling for Uncertainty

	Tobin’s Q					
	Patent			Wiki		
	(1)	(2)	(3)	(4)	(5)	(6)
Predicted Workstyle Change in Industry \times Log Age	-0.159*** (0.043)	-0.158*** (0.043)	-0.153** (0.060)	-0.148*** (0.045)	-0.150*** (0.045)	-0.150** (0.059)
Predicted Workstyle Change in Industry \times Log Size	0.045* (0.020)	0.043* (0.019)	0.027 (0.016)	0.047** (0.018)	0.045** (0.018)	0.029* (0.016)
Predicted Workstyle Change in Industry	0.290* (0.175)	0.304* (0.174)	0.398* (0.226)	0.199 (0.192)	0.215 (0.194)	0.338 (0.272)
Uncertainty \times Log Age	-0.042 (0.138)	0.002** (0.001)	-0.150 (0.156)	-0.048 (0.137)	0.002** (0.001)	-0.143 (0.155)
Uncertainty \times Log Size	-0.100 (0.072)	-0.001** (0.001)	-0.146*** (0.047)	-0.100 (0.072)	-0.001** (0.001)	-0.147*** (0.047)
Uncertainty	0.990** (0.579)	0.004 (0.005)	0.606 (0.589)	1.012** (0.572)	0.004 (0.005)	0.576 (0.590)
Log Age	-0.245*** (0.064)	-0.507*** (0.128)	-0.144 (0.109)	-0.267*** (0.069)	-0.552*** (0.137)	-0.167 (0.107)
Log Size	-0.173*** (0.034)	-0.040 (0.064)	-0.107** (0.045)	-0.168*** (0.034)	-0.029 (0.064)	-0.106** (0.045)
Uncertainty Measure	WUI	EPU	Return Vol	WUI	EPU	Return Vol
Year FE	X	X	X	X	X	X
Observations	27,978	27,978	16,811	27,978	27,978	16,811
Within R^2	0.077	0.077	0.095	0.077	0.076	0.093

Notes: This table presents regressions at the firm-year level in Equation (41). The independent variables include the predicted workstyle change in industry n over the next 5 years due to year t technologies ($|\overline{dws}|_{n,t,5}^e$), and its interaction with firm age and firm size (log book assets in year t). We standardize $|\overline{dws}|_{n,t,5}^e$ to unit variance and zero median. Columns (1) to (3) use the predicted workstyle change based on breakthrough patents as in Tables 2 to 4, and columns (4) to (6) use the predicted workstyle change based on Wikipedia titles as in Tables IA3 to IA5. All columns additionally control for an uncertainty measure and its interaction with firm age and firm size, where uncertainty is proxied by the World Uncertainty Index (WUI) in columns (1) and (4), economic policy uncertainty (EPU) from Baker, Bloom, and Davis (2016) in columns (2) and (5), and firm-level stock return volatility (Return Vol) in columns (3) and (6). We include year fixed effects, and bootstrap standard errors. In each bootstrap draw, we resample industries with replacement, and use the corresponding bootstrap weights to reestimate the technology-implied employment change by industry-occupation-year and recalculate the predicted workstyle change, which is then used to rerun the final regression of interest in the bootstrap sample. The coefficient is the bootstrap median. Asterisks denote significance levels (***=1%, **=5%, *=10%). Number of observations and within R^2 (without fixed effects) are based on the OLS regression. Sample years are 2003 to 2016.

Table IA8 – Sales Growth among Compustat Firms, Controlling for Uncertainty

	Sales Growth (5-Year DHS Rate)					
	Patent			Wiki		
	(1)	(2)	(3)	(4)	(5)	(6)
Predicted Workstyle Change in Industry \times Log Age	-0.026** (0.009)	-0.026** (0.009)	-0.021** (0.008)	-0.026** (0.009)	-0.026** (0.009)	-0.021** (0.008)
Predicted Workstyle Change in Industry \times Log Size	0.010* (0.005)	0.011* (0.005)	0.009 (0.005)	0.012** (0.006)	0.012** (0.005)	0.011 (0.006)
Predicted Workstyle Change in Industry	0.020 (0.032)	0.018 (0.031)	0.009 (0.035)	0.008 (0.037)	0.009 (0.035)	-0.009 (0.042)
Uncertainty \times Log Age	0.026 (0.024)	0.000*** (0.000)	0.072*** (0.028)	0.025 (0.024)	0.001*** (0.000)	0.073*** (0.028)
Uncertainty \times Log Size	0.007 (0.011)	0.000 (0.000)	-0.047** (0.015)	0.007 (0.011)	0.000 (0.000)	-0.047** (0.015)
Uncertainty	-0.148 (0.097)	-0.003*** (0.001)	-0.255*** (0.103)	-0.142 (0.097)	-0.003*** (0.001)	-0.258*** (0.102)
Log Age	-0.103*** (0.018)	-0.167*** (0.032)	-0.126*** (0.020)	-0.105*** (0.018)	-0.173*** (0.033)	-0.129*** (0.020)
Log Size	-0.026** (0.012)	-0.046** (0.017)	-0.024* (0.012)	-0.025** (0.012)	-0.043** (0.017)	-0.023* (0.012)
Uncertainty Measure	WUI	EPU	Return Vol	WUI	EPU	Return Vol
Year FE	X	X	X	X	X	X
Observations	27,978	27,978	16,811	27,978	27,978	16,811
Within R^2	0.059	0.061	0.070	0.060	0.061	0.071

Notes: This table presents regressions at the firm-year level in Equation (42), where the outcome variable $\Delta Y_{i,t \rightarrow t+5}$ is sales growth between year t and $t + 5$ calculated following Davis, Haltiwanger, and Schuh (1992) as $\Delta Y_{i,t \rightarrow t+5} = (Y_{i,t+5} - Y_{i,t}) / (0.5 \times Y_{i,t+5} + 0.5 \times Y_{i,t})$. The independent variables include the predicted workstyle change in industry n over the next 5 years due to year t technologies ($|d\bar{w}s|_{n,t,5}^e$), and its interaction with firm age and firm size (log sales in year t). We standardize $|d\bar{w}s|_{n,t,5}^e$ to unit variance and zero median. Columns (1) to (3) use the predicted workstyle change based on breakthrough patents as in Tables 2 to 4, and columns (4) to (6) use the predicted workstyle change based on Wikipedia titles as in Tables IA3 to IA5. All columns additionally control for an uncertainty measure and its interaction with firm age and firm size, where uncertainty is proxied by the World Uncertainty Index (WUI) in columns (1) and (4), economic policy uncertainty (EPU) from Baker, Bloom, and Davis (2016) in columns (2) and (5), and firm-level stock return volatility (Return Vol) in columns (3) and (6). We include year fixed effects, and bootstrap standard errors. In each bootstrap draw, we resample industries with replacement, and use the corresponding bootstrap weights to reestimate the technology-implied employment change by industry-occupation-year and recalculate the predicted workstyle change, which is then used to rerun the final regression of interest in the bootstrap sample. The coefficient is the bootstrap median. Asterisks denote significance levels (***=1%, **=5%, *=10%). Number of observations and within R^2 (without fixed effects) are based on the OLS regression. Sample years are 2003 to 2016.

Table IA9 – Employment Growth among Compustat Firms, Controlling for Uncertainty

	Employment Growth (5-Year DHS Rate)					
	Patent			Wiki		
	(1)	(2)	(3)	(4)	(5)	(6)
Predicted Workstyle Change in Industry \times Log Age	-0.022** (0.007)	-0.022** (0.007)	-0.019** (0.008)	-0.022** (0.007)	-0.022** (0.007)	-0.021** (0.008)
Predicted Workstyle Change in Industry \times Log Size	0.009* (0.004)	0.009* (0.004)	0.007 (0.004)	0.009** (0.003)	0.009** (0.003)	0.007 (0.004)
Predicted Workstyle Change in Industry	0.078** (0.026)	0.077** (0.026)	0.075*** (0.025)	0.076** (0.027)	0.077** (0.026)	0.080*** (0.025)
Uncertainty \times Log Age	0.011 (0.025)	0.000 (0.000)	0.053** (0.030)	0.010 (0.025)	0.000 (0.000)	0.054** (0.030)
Uncertainty \times Log Size	-0.015 (0.009)	0.000 (0.000)	-0.015 (0.013)	-0.014 (0.009)	0.000 (0.000)	-0.015 (0.013)
Uncertainty	-0.052 (0.098)	-0.001 (0.001)	-0.374*** (0.124)	-0.050 (0.098)	-0.001 (0.001)	-0.376*** (0.125)
Log Age	-0.093*** (0.015)	-0.139*** (0.036)	-0.110*** (0.022)	-0.095*** (0.015)	-0.144*** (0.037)	-0.111*** (0.022)
Log Size	-0.006 (0.006)	-0.017 (0.014)	-0.015 (0.010)	-0.005 (0.006)	-0.015 (0.014)	-0.014 (0.009)
Uncertainty Measure	WUI	EPU	Return Vol	WUI	EPU	Return Vol
Year FE	X	X	X	X	X	X
Observations	27,978	27,978	16,811	27,978	27,978	16,811
Within R^2	0.040	0.041	0.039	0.039	0.041	0.039

Notes: This table presents regressions at the firm-year level in Equation (42), where the outcome variable $\Delta Y_{i,t \rightarrow t+5}$ is employment growth between year t and $t + 5$ calculated following [Davis, Haltiwanger, and Schuh \(1992\)](#) as $\Delta Y_{i,t \rightarrow t+5} = (Y_{i,t+5} - Y_{i,t}) / (0.5 \times Y_{i,t+5} + 0.5 \times Y_{i,t})$. The independent variables include the predicted workstyle change in industry n over the next 5 years due to year t technologies ($|d\bar{w}s|_{n,t,5}^e$), and its interaction with firm age and firm size (log employment in year t). We standardize $|d\bar{w}s|_{n,t,5}^e$ to unit variance and zero median. Columns (1) to (3) use the predicted workstyle change based on breakthrough patents as in [Tables 2 to 4](#), and columns (4) to (6) use the predicted workstyle change based on Wikipedia titles as in [Tables IA3 to IA5](#). All columns additionally control for an uncertainty measure and its interaction with firm age and firm size, where uncertainty is proxied by the World Uncertainty Index (WUI) in columns (1) and (4), economic policy uncertainty (EPU) from [Baker, Bloom, and Davis \(2016\)](#) in columns (2) and (5), and firm-level stock return volatility (Return Vol) in columns (3) and (6). We include year fixed effects, and bootstrap standard errors. In each bootstrap draw, we resample industries with replacement, and use the corresponding bootstrap weights to reestimate the technology-implied employment change by industry-occupation-year and recalculate the predicted workstyle change, which is then used to rerun the final regression of interest in the bootstrap sample. The coefficient is the bootstrap median. Asterisks denote significance levels (***=1%, **=5%, *=10%). Number of observations and within R^2 (without fixed effects) are based on the OLS regression. Sample years are 2003 to 2016.

Table IA10 – Tobin’s Q among Compustat Firms, Controlling for Profitability

	Tobin’s Q					
	Patent			Wiki		
	(1)	(2)	(3)	(4)	(5)	(6)
Predicted Workstyle Change in Industry \times Log Age	-0.122** (0.040)	-0.134*** (0.039)	-0.160*** (0.044)	-0.114** (0.039)	-0.127** (0.041)	-0.148*** (0.046)
Predicted Workstyle Change in Industry \times Log Size	0.046** (0.016)	0.046* (0.019)	0.037 (0.020)	0.048** (0.016)	0.047** (0.018)	0.040* (0.018)
Predicted Workstyle Change in Industry \times Profitability	-0.197 (0.304)	-0.032 (0.038)	0.070 (0.063)	-0.268 (0.297)	-0.037 (0.041)	0.042 (0.062)
Predicted Workstyle Change in Industry	0.167 (0.160)	0.209 (0.184)	0.350* (0.190)	0.104 (0.175)	0.142 (0.204)	0.242 (0.220)
Log Age	-0.187*** (0.061)	-0.232*** (0.073)	-0.275*** (0.072)	-0.206*** (0.064)	-0.252*** (0.077)	-0.300*** (0.076)
Log Size	-0.113*** (0.024)	-0.164*** (0.034)	-0.194*** (0.036)	-0.108*** (0.024)	-0.159*** (0.035)	-0.191*** (0.037)
Profitability	-1.035*** (0.255)	-0.049*** (0.068)	0.305*** (0.057)	-1.072*** (0.233)	-0.059*** (0.063)	0.323*** (0.061)
Profitability Measure	NI/Assets	NI/Sales	E/P	NI/Assets	NI/Sales	E/P
Year FE	X	X	X	X	X	X
Observations	27,084	27,204	27,007	27,084	27,204	27,007
Within R^2	0.129	0.092	0.081	0.127	0.092	0.080

Notes: This table presents regressions at the firm-year level in Equation (41). The independent variables include the predicted workstyle change in industry n over the next 5 years due to year t technologies ($|d\bar{w}s|_{n,t,5}^e$), and its interaction with firm age and firm size (log book assets in year t). We standardize $|d\bar{w}s|_{n,t,5}^e$ to unit variance and zero median. Columns (1) to (3) use the predicted workstyle change based on breakthrough patents as in Tables 2 to 4, and columns (4) to (6) use the predicted workstyle change based on Wikipedia titles as in Tables IA3 to IA5. All columns additionally control for the interaction of $|d\bar{w}s|_{n,t,5}^e$ and a measure of current profitability to capture concerns of cannibalization, which uses net income over assets (NI/Assets) in columns (1) and (4), net income over sales (NI/Sales) in columns (2) and (5), and the earnings-to-price ratio (E/P) in columns (3) and (6). We include year fixed effects, and bootstrap standard errors. In each bootstrap draw, we resample industries with replacement, and use the corresponding bootstrap weights to reestimate the technology-implied employment change by industry-occupation-year and recalculate the predicted workstyle change, which is then used to rerun the final regression of interest in the bootstrap sample. The coefficient is the bootstrap median. Asterisks denote significance levels (***=1%, **=5%, *=10%). Number of observations and within R^2 (without fixed effects) are based on the OLS regression. Sample years are 2003 to 2016.

Table IA11 – Sales Growth among Compustat Firms, Controlling for Profitability

	Sales Growth (5-Year DHS Rate)					
	Patent			Wiki		
	(1)	(2)	(3)	(4)	(5)	(6)
Predicted Workstyle Change in Industry \times Log Age	-0.024** (0.008)	-0.024** (0.008)	-0.023* (0.009)	-0.025** (0.008)	-0.024** (0.008)	-0.023* (0.009)
Predicted Workstyle Change in Industry \times Log Size	0.010* (0.005)	0.009* (0.004)	0.010 (0.005)	0.011* (0.005)	0.008* (0.004)	0.011* (0.006)
Predicted Workstyle Change in Industry \times Profitability	0.019 (0.037)	0.003 (0.006)	0.002 (0.012)	0.023 (0.040)	0.003 (0.007)	0.000 (0.013)
Predicted Workstyle Change in Industry	0.016 (0.032)	0.031 (0.028)	0.013 (0.034)	0.011 (0.035)	0.031 (0.028)	0.000 (0.040)
Log Age	-0.092*** (0.017)	-0.099*** (0.017)	-0.098*** (0.017)	-0.095*** (0.017)	-0.102*** (0.017)	-0.101*** (0.017)
Log Size	-0.018* (0.010)	-0.006 (0.007)	-0.025** (0.013)	-0.017 (0.010)	-0.005 (0.007)	-0.025** (0.012)
Profitability	-0.108** (0.038)	-0.023*** (0.006)	0.092*** (0.023)	-0.100** (0.034)	-0.022*** (0.006)	0.093*** (0.023)
Profitability Measure	NI/Assets	NI/Sales	E/P	NI/Assets	NI/Sales	E/P
Year FE	X	X	X	X	X	X
Observations	27,084	27,204	27,007	27,084	27,204	27,007
Within R^2	0.060	0.056	0.058	0.060	0.056	0.059

Notes: This table presents regressions at the firm-year level in Equation (42), where the outcome variable $\Delta Y_{i,t \rightarrow t+5}$ is sales growth between year t and $t + 5$ calculated following Davis, Haltiwanger, and Schuh (1992) as $\Delta Y_{i,t \rightarrow t+5} = (Y_{i,t+5} - Y_{i,t}) / (0.5 \times Y_{i,t+5} + 0.5 \times Y_{i,t})$. The independent variables include the predicted workstyle change in industry n over the next 5 years due to year t technologies ($|d\bar{w}s|_{n,t,5}^e$), and its interaction with firm age and firm size (log sales in year t). We standardize $|d\bar{w}s|_{n,t,5}^e$ to unit variance and zero median. Columns (1) to (3) use the predicted workstyle change based on breakthrough patents as in Tables 2 to 4, and columns (4) to (6) use the predicted workstyle change based on Wikipedia titles as in Tables IA3 to IA5. All columns additionally control for the interaction of $|d\bar{w}s|_{n,t,5}^e$ and a measure of current profitability to capture concerns of cannibalization, which uses net income over assets (NI/Assets) in columns (1) and (4), net income over sales (NI/Sales) in columns (2) and (5), and the earnings-to-price ratio (E/P) in columns (3) and (6). We include year fixed effects, and bootstrap standard errors. In each bootstrap draw, we resample industries with replacement, and use the corresponding bootstrap weights to reestimate the technology-implied employment change by industry-occupation-year and recalculate the predicted workstyle change, which is then used to rerun the final regression of interest in the bootstrap sample. The coefficient is the bootstrap median. Asterisks denote significance levels (***=1%, **=5%, *=10%). Number of observations and within R^2 (without fixed effects) are based on the OLS regression. Sample years are 2003 to 2016.

Table IA12 – Employment Growth among Compustat Firms, Controlling for Profitability

	Employment Growth (5-Year DHS Rate)					
	Patent			Wiki		
	(1)	(2)	(3)	(4)	(5)	(6)
Predicted Workstyle Change in Industry \times Log Age	-0.024** (0.007)	-0.023** (0.007)	-0.019** (0.007)	-0.025** (0.007)	-0.023** (0.007)	-0.020** (0.007)
Predicted Workstyle Change in Industry \times Log Size	0.008* (0.003)	0.009** (0.003)	0.008 (0.004)	0.007* (0.003)	0.009** (0.003)	0.008* (0.004)
Predicted Workstyle Change in Industry \times Profitability	0.044 (0.034)	0.000 (0.004)	0.005 (0.019)	0.062* (0.040)	0.001 (0.005)	-0.003 (0.017)
Predicted Workstyle Change in Industry	0.085*** (0.025)	0.080** (0.025)	0.070** (0.025)	0.087** (0.025)	0.079** (0.025)	0.068** (0.027)
Log Age	-0.089*** (0.014)	-0.090*** (0.014)	-0.094*** (0.014)	-0.092*** (0.014)	-0.092*** (0.014)	-0.096*** (0.014)
Log Size	-0.010* (0.006)	-0.010* (0.006)	-0.012* (0.006)	-0.009 (0.006)	-0.009 (0.006)	-0.011 (0.007)
Profitability	0.015 (0.058)	0.003* (0.008)	0.154*** (0.024)	0.023 (0.051)	0.003** (0.008)	0.159*** (0.024)
Profitability Measure	NI/Assets	NI/Sales	E/P	NI/Assets	NI/Sales	E/P
Year FE	X	X	X	X	X	X
Observations	27,084	27,204	27,007	27,084	27,204	27,007
Within R^2	0.040	0.040	0.049	0.040	0.039	0.048

Notes: This table presents regressions at the firm-year level in Equation (42), where the outcome variable $\Delta Y_{i,t \rightarrow t+5}$ is employment growth between year t and $t + 5$ calculated following Davis, Haltiwanger, and Schuh (1992) as $\Delta Y_{i,t \rightarrow t+5} = (Y_{i,t+5} - Y_{i,t}) / (0.5 \times Y_{i,t+5} + 0.5 \times Y_{i,t})$. The independent variables include the predicted workstyle change in industry n over the next 5 years due to year t technologies ($|d\bar{w}s|_{n,t,5}^e$), and its interaction with firm age and firm size (log employment in year t). We standardize $|d\bar{w}s|_{n,t,5}^e$ to unit variance and zero median. Columns (1) to (3) use the predicted workstyle change based on breakthrough patents as in Tables 2 to 4, and columns (4) to (6) use the predicted workstyle change based on Wikipedia titles as in Tables IA3 to IA5. All columns additionally control for the interaction of $|d\bar{w}s|_{n,t,5}^e$ and a measure of current profitability to capture concerns of cannibalization, which uses net income over assets (NI/Assets) in columns (1) and (4), net income over sales (NI/Sales) in columns (2) and (5), and the earnings-to-price ratio (E/P) in columns (3) and (6). We include year fixed effects, and bootstrap standard errors. In each bootstrap draw, we resample industries with replacement, and use the corresponding bootstrap weights to reestimate the technology-implied employment change by industry-occupation-year and recalculate the predicted workstyle change, which is then used to rerun the final regression of interest in the bootstrap sample. The coefficient is the bootstrap median. Asterisks denote significance levels (***=1%, **=5%, *=10%). Number of observations and within R^2 (without fixed effects) are based on the OLS regression. Sample years are 2003 to 2016.

IA2 Processing Wikipedia Titles

This appendix describes in detail how we obtain and process Wikipedia titles on technological inventions used in Section II.

IA2.1 Wikipedia Titles on Technological Inventions

We extract English Wikipedia titles from the [Wikimedia English Wikipedia Dumps](#) and filter them in two steps to identify titles related to technological inventions.

Step 1 We first do batch filtering to obtain a broad set of technology-related titles using GPT-4o-mini. The LLM is instructed to identify Wikipedia titles corresponding to technological inventions. Wikipedia titles are processed in batches of 50, and the LLM returns only those classified as technological inventions.

Instructions:

- You are tasked with filtering a list of Wikipedia article titles to include only those of technological inventions. A technological invention is a novel device, method, process, or system that applies scientific or engineering principles to solve a problem or improve efficiency. These inventions typically introduce new functionalities, enhance existing technologies, or create entirely new categories of tools or systems.
- Focus ONLY on technological inventions, which are devices, systems, and components with novel technological advancements, such as 'wheels', 'artificial intelligence', 'electric vehicle', 'mobile phone', 'Watt steam engine', etc.
- EXCLUDE names of people, movies, books, songs, food, animals, places, companies, organizations, social movements, slogans, cultural phrases, and any non-technology items.
- EXCLUDE specific models, product names, versions of weapons and vehicles, unless they represent a significant technological advancement (e.g., 'iPhone' is a technological invention, but 'iPhone 13' is not. Airplane is a technological invention, but 'Boeing 747' is not).
- Do NOT include scientific phenomena unless they describe a concrete technological invention or device.
- If multiple titles are extremely similar (such as variations in capitalization, punctuation, or wording), ONLY include ONE representative title.

For each title, first assess if it represents a technological invention. Then, return ONLY the valid titles (keeping the original format of input, including the '_' character).

Respond in a JSON array of strings like this:

```
['Title 1', 'Title 2', 'Title 3']
```

Titles:

{titles_list}

Step 2 We validate each Wikipedia title from Step 1 using its Wikipedia summary to obtain the final set of technological inventions. For each title, we use the Wikipedia API library to retrieve the corresponding English Wikipedia summary. We then supply both the title and the summary to Deepseek-Reasoner model, which is prompted to determine whether the article indeed describes a technological invention, following the same definition used in Step 1. The LLM's responses indicate whether the Wikipedia title should be retained.

Instructions:

- You are tasked with determining whether a given Wikipedia article is about a technological invention based on the title and summary provided. A technological invention is a novel device, method, process, or system that applies scientific or engineering principles to solve a problem or improve efficiency. These inventions typically introduce new functionalities, enhance existing technologies, or create entirely new categories of tools or systems.

- EXCLUDE specific models, product names, versions of weapons and vehicles, unless they represent a significant technological advancement (e.g., 'iPhone' is a technological invention, but 'iPhone 13' is not. Airplane is a technological invention, but 'Boeing 747' is not). Example of titles that should be EXCLUDED: 'AJS_Model_16', 'AMD_2900', 'A340_(aircraft)', 'A.T_Mine_E.P._Mark_II'.

Wikipedia Article Title: {title}

Wikipedia Article Summary:

\\''\\''\\''{summary}\\''\\''\\''

Based on the above summary, determine if this article is about a technological invention.

Respond ONLY with a JSON object in the following format:

{{'is_valid': true}}

or

{{'is_valid': false}}

IA2.2 Technology Information

We use the following prompt to summarize basic information about the technological inventions with Gemini-2.0-Flash, including the time and location of the invention, as well as the type of inventor (i.e., private company, individual inventor, government, or university/non-profit).

Instructions:

You are an expert in the history of technologies. Your task is to answer questions about when, where, who invented the technology, and which category of application it best fits into, based on the given information about the technology and your domain knowledge.

Provide the following information along with a brief reason (less than 20 words) for each:

- 1) When the technology first became workable. If the exact year is unavailable or uncertain, provide the closest approximation. Be as accurate as possible. Include 'BCE' in your response if the date is before 0 CE/very ancient.
- 2) Where it was first created (country where the first workable version was created).
- 3) Who invented it (choose one of these categories by the type of inventor):
 - 'Private Company': For-profit company/companies or people working for these companies developed the technology (e.g. Apple, Google);
 - 'University/Non-Profit': Non-profit institution or organization developed it, such as universities, non-profit institutions, open-source projects (e.g. UC Berkeley, SRI, Mozilla);
 - 'Government Project': Developed via government-funded or military project (e.g. NASA, the Apollo program, the Manhattan Project);
 - 'Individual Inventor': Developed by an individual or a small team of inventors unaffiliated with large institutions or companies (e.g. Nikola Tesla, Philo Farnsworth);

Note: Some inventions involve collaboration. When multiple parties are involved, choose the category that best reflects the majority of developers or the most responsible entity. Aim for the most accurate classification based on available evidence. When inventors are not explicitly mentioned, use your best judgment to classify the technology based on the context and purpose of the invention.

For example:

- A collaboration between private companies is still classified as 'Private Company'.
- A small team of unaffiliated inventors is still 'Individual Inventor'.
- If an individual works under a government-funded initiative, classify it as 'Government Project'.
- If it was published by an industry association, it should be either 'Private Company' or 'University/Non-Profit' depending on the inventor and purpose of the technology. For example, SAE (Society of Automotive Engineers) is mostly ran by engineers from private companies, so it should be classified as 'Private Company'. W3C (World Wide Web Consortium) is ran by university researchers, so it should be classified as 'University/Non-Profit'.
- If it was designed for commercial use or production efficiency, it's likely 'Private Company'.
- If it served military, national defense, or public interest, it's likely 'Government Project'.
- If it was developed for academic research or knowledge advancement, it's likely 'University/Non-Profit'.
- If it was created by individuals not working for large institutions, it's likely 'Individual Inventor'.

4) The 'domain_of_application' (choose exactly one from the list below):

['Agriculture, Forestry, Fishing and Hunting', 'Mining, Quarrying, and Oil and Gas Extraction', 'Utilities', 'Construction', 'Manufacturing', 'Wholesale Trade', 'Retail Trade', 'Transportation and Warehousing', 'Information', 'Finance and Insurance', 'Real Estate and Rental and Leasing', 'Professional, Scientific, and Technical Services', 'Health Care and Social Assistance', 'Military and Defense', 'Basic Research']

Importantly: Use only the listed categories above. Do not invent new categories.

Technology Name: {title}

Wikipedia Page:

```
\''\''\''{wiki_text[:5000]}\''\''\''\''
```

Return ONLY a JSON object in the following format:

```
{{
  'technology': 'Title',
  'when': year or 'Not Available',
  'when_reason': 'Brief reason for chosen year',
  'where': country name or 'Not Available',
  'where_reason': 'Brief reason for chosen country',
  'who': one of the 4 listed categories,
  'who_reason': 'Brief reason for chosen category',
  'domain_of_application': one of the 15 listed categories,
  'domain_reason': 'Brief reason for chosen domain'
}}
```

Example Output #1:

```
{{
  'technology': 'iPhone',
  'when': '2007',
  'when_reason': 'Apple released the first iPhone in 2007.',
  'where': 'USA',
  'where_reason': 'Apple developed the iPhone in the United States.',
  'who': 'Private Company',
  'who_reason': 'Developed by Apple, which is a for profit company.',
  'domain_of_application': 'Information',
  'domain_reason': 'Primarily used for communication and computing.'
}}
```

Example Output #2:

```
{{
  'technology': 'CRISPR',
  'when': '2012',
  'when_reason': 'First practical demonstration of gene editing was in
2012.',
}}
```

```
'where': 'USA',
'where_reason': 'Initial research conducted at UC Berkeley.',
'who': 'University Lab',
'who_reason': 'Developed at UC Berkeley by academic researchers.',
'domain_of_application': 'Health Care and Social Assistance',
'domain_reason': 'Primarily applied to medicine and biotechnology.'
}}
```

Use this JSON schema:

```
{
  'technology': 'str',
  'when': 'str',
  'when_reason': 'str',
  'where': 'str',
  'where_reason': 'str',
  'who': 'str',
  'who_reason': 'str',
  'domain_of_application': 'str',
  'domain_reason': 'str'
}
```

IA2.3 Technology Implementation and Commercialization

We use the following prompt to ask web-enabled Gemini-2.0-Flash whether the company that was most successful in its initial implementation and commercialization was a young firm, or an old incumbent.

You are a historian of technology and innovation.
You must use web search to verify the answer.

Here is a technological invention.
Your task is to determine whether the given technology was ever commercialized,
identify the company that was most successful in its initial implementation and commercialization,
and classify whether that company was a young firm or an old firm at the time of the commercialization.

INSTRUCTIONS

Step 1: Identify the implementing company if the technology was ever commercialized

- Return the name of the company that was most successful in the initial implementation and commercialization of the technology.
- Always use the name of the parent company and do not use the name of the division or subsidiary.
- If multiple companies were involved, return the primary company most widely credited with initial commercialization.
- If the technology has never been commercialized, return "Never Commercialized".

Step 2: Classify firm type

Determine whether the implementing company was young or old at the time of the commercialization, based on the age of the parent company:

- ****1**** → Young firm (founded less than 10 years before commercialization)
- ****1**** → Old firm (existed long before the invention)
- ****0**** → Not sure or ambiguous
- ****2**** → Non-commercialized (never commercialized)

Step 3: Provide short reasoning and approximate period

INPUT

Technology: {title}

Wikipedia summary:

\{\summary_text\}

OUTPUT FORMAT (JSON ONLY)

Return only valid JSON:

```
{  
  "technology": "{title}",  
  "implementing_firm_name": "Name of company or 'Never Commercialized'",  
  "firm_type": 1 or -1 or 0 or 2,  
  "firm_type_reason": "Short explanation (<30 words)",  
  "year_or_period": "Approximate year/decade of first commercialization if  
  available, else 'Unknown'"  
}
```

Do not include any explanation outside the JSON.

EXAMPLES

Dropbox

```
{  
  "technology": "Dropbox",  
  "implementing_firm_name": "Dropbox",  
  "firm_type": 1,  
  "firm_type_reason": "Founded and launched Dropbox in 2007.",  
  "year_or_period": "2007"  
}
```

iPhone

```
{  
  "technology": "iPhone",  
  "implementing_firm_name": "Apple",  
  "firm_type": -1,  
  "firm_type_reason": "Apple founded in 1976; launched iPhone in 2007.",  
  "year_or_period": "2007"  
}
```

Ethernet

```
{  
  "technology": "Ethernet",  
  "implementing_firm_name": "3Com",  
  "firm_type": 1,  
  "firm_type_reason": "3Com founded in 1979; commercialized Ethernet hardware  
in the early 1980s.",  
  "year_or_period": "1980s"  
}
```

Asteroid Impact Avoidance

```
{  
  "technology": "Asteroid Impact Avoidance",  
  "implementing_firm_name": "Never Commercialized",  
  "firm_type": 2,  
  "firm_type_reason": "Remained research; no market commercialization.",  
  "year_or_period": "Unknown"  
}
```

IA3 Processing Employee Reviews

We collect information through employee reviews from Revelio ([Revelio Labs, 2025](#)). We process the reviews using the prompts below.

IA3.1 Prevalence of Rules

Step 1 We process all employee reviews with Gemini-2.0-Flash to screen for those that discuss rules.

You are analyzing employee reviews of companies. Your task is to determine whether this review discusses rules, protocols, procedures, policies, or norms that the employee's company has. Rules include formal policies and processes as well as strong unwritten norms.

- Ignore comments about pay, culture, or leadership that do not explicitly mention rules or norms.
- We are looking for reviews about rules in the reviewer's employer company, not rules imposed by the regulators, suppliers, or clients of the reviewer's employer company.

Here is a piece of employee review:

```
''{text}''
```

Does the review discuss rules, protocols, procedures, policies, or norms that the employee's company has?

- Yes (return 1)
- No (return -1)
- Unclear (return 0)

Respond ONLY with a JSON object in the following format:

```
{{'flag': 0, 'confidence': 0.5}}
```

Step 2 We process reviews labeled as 1 or 0 from Step 1 with Gemini-2.0-Flash to determine whether they mention the presence or the lack of rules.

Here is a piece of employee review for a company:

```
''{text}''
```

Please determine which of the following best describes this review:

- This review discusses the presence of rules and processes in the company (return 1)
- This review discusses the lack of rules and processes in the company (return -1)
- This review does not discuss rules and processes in the company (return 0)

Respond ONLY with a JSON object in the following format:

```
{  
  'label': 0,  
  'confidence': 0.5,  
  'snippets': (if the label is not 0, report the snippets from the review  
               that was most important for your determination of the label)  
}
```

Rule index We use the results from Step 2 in the numerator of the rule index, and calculate the intensity of rules in a GVKEY-year as:

$$\text{Rule Index} = \frac{(\# \text{ of reviews mentioning the presence of rules} - \# \text{ of reviews mentioning the lack of rules})}{\# \text{ of total reviews}}$$

We map the Revelio company ID RCID to GVKEY in a given year using the parent-subsidary bridge provided by Revelio to us directly.

Examples of employee review discussions

- Presence of rules:
 - “It’s a very structured place with a lot of rules, hierarchy, etc.”
 - “Stick to the rules and not flexible.”
 - “Lots of rules and guidelines.”
 - “A stable and well structured corporation that follows the rules.”
 - “Company follows strict rules in all type of scenario.”
- Absence of rules:
 - “Unstructured work environment. Lack of leadership. Lack of accountability - Job Functions are not very well defined and will change constantly.”

- "Lack of, or nonexistent enforcement of policy results in extremely inconsistent standards."
- "Make many changes at the last minute...schedules, resource guide, rules. etc."
- "Need to improve culture and standardize HR rules."
- "Rules of engagement and process were not always clear."

IA3.2 Too Many Meetings

Step 1 We process all employee reviews with Gemini-2.0-Flash to screen for those that discuss meetings, committees, and layers of approval.

You are analyzing employee reviews of companies. Your task is to determine whether this review discusses meetings, committees, and layers of approval.

- We are looking for reviews about such activities in the reviewer's employer company, not those imposed by the regulators, suppliers, or clients of the reviewer's employer company or social meetings and parties.

- Do not include personal, non-work related activities, such as smoking, drinking, eating, or playing games.

- Do not include general discussions about office politics, slow promotion, or bad management.

Here is a piece of employee review:

```
''{text}''
```

Does the review discuss meetings, committees, and layers of approval at the employee's company?

- Yes (return 1)

- No (return -1)

- Unclear (return 0)

Respond ONLY with a JSON object in the following format:

```
{{'flag': 0, 'confidence': 0.5}}
```

Step 2 We process reviews labeled as 1 or 0 from Step 1 with Gemini-2.0-Flash to determine whether they complain that the company has too many meetings, committees, and layers of approval or praise the company for not having too many.

Here is a piece of employee review for a company:

```
''{text}''
```

Please determine which of the following best describes this review:

- This review complains about too many meetings, committees, and layers of approval that waste time. (return 1)
- This review praises the company for being efficient and not having too many meetings, committees, and layers of approval that waste time. (return -1)
- It is difficult to tell (return 0)

Respond ONLY with a JSON object in the following format:

```
{  
  'label': 0,  
  'confidence': 0.5,  
  'snippets': (if the label is not 0, report the snippets from the review  
               that was most important for your determination of the label)  
}
```

Meeting index We use the results from Step 2 in the numerator of the meeting index, and calculate the intensity of meetings that waste time in a GVKEY-year as:

$$\text{Meeting Index} = \frac{(\# \text{ of reviews complaining too many} - \# \text{ of reviews praising not too many})}{\# \text{ of total reviews}}.$$

We map Revelio company ID RCID to GVKEY in a given year using parent-subsiary bridge provided by Revelio to us directly.

Examples of employee review discussions

- Too many meetings:
 - “Way too many meetings. Took forever to make changes and/or decisions.”
 - “More meetings and games than work.”
 - “Too many internal mandatory calls which takes away time best served somewhere else.”
 - “A lot of meaningless meetings.”
 - “Too much internal fighting, more competition with folks inside rather than outside.”
- Not too many meetings:

- “Meeting-light: this was a HUGE unlock for me. My previous job had me back to back in meetings from 9a-5p so the only time I had to get work done was AFTER work.”
- “Good work environment, not too many meetings during the day.”
- “The weekly meetings were not excessive.”
- “There are lots of resources and projects going on, and very few fiefdoms or silos to get in the way of you contributing.”
- “Not being micro-managed, feeling challenged, being a part of something bigger than just a job, but a disruption of an industry, with camaraderie along the way are things I value.”

IA4 Measuring Similarity between Technologies and Occupation Tasks

This appendix provides additional details on the computation of textual similarity between new technologies and occupations’ routine and non-routine task descriptions, as used in Sections IV.B and VI.A.1. The overall approach follows Kogan et al. (2024): we represent each technology document and each occupation’s task descriptions as vectors in embedding space, and compute cosine similarity between them. Section IA4.1 describes the text preprocessing and embedding construction for the patent-based measures. Section IA4.2 describes the analogous procedure for Wikipedia titles, including the extraction of technical keywords used to construct Wikipedia title embeddings.

IA4.1 Similarity between Breakthrough Patents and Occupations’ Task Descriptions

This subsection describes the text preprocessing steps applied to patent text and O*NET occupation task descriptions before computing the document embeddings \mathcal{X}_p and \mathcal{X}_j^r in Equation (34) and the cosine similarity $s_{p,j}^r$ in Equation (35) in Section IV.B.

Patent text We obtain patent text from PatentsView. For each utility patent, we combine the abstract, detailed description, and all claims into a single text document. We then apply the following preprocessing steps:

1. Remove non-alphabetic characters.
2. Convert to lowercase.
3. Tokenize using NLTK.
4. Part-of-speech (POS) tag using NLTK.

5. Keep only nouns and verbs.
6. Remove stopwords (using a combined stopword list from multiple sources, including NLTK, spaCy, Snowball, SMART, Lextek, ONIX, and patent-specific stopwords from [Kogan et al. \(2024\)](#) such as “abstract,” “claim,” “invention,” “patent,” “united,” and “states”).
7. Lemmatize using WordNet.

Occupation task text We apply the same preprocessing pipeline to O*NET occupation task descriptions. As described in Section [IV.B](#), for each occupation, tasks are classified as routine or non-routine using LLM (GPT-4o). We then concatenate all routine task descriptions into one document and all non-routine task descriptions into another, producing two text documents per occupation.

Embeddings As described in Section [IV.B](#), each document X_i —whether a patent or an occupation’s combined routine (non-routine) task descriptions—is represented as:

$$X_i = \sum_k q_{i,k} x_k,$$

where x_k is the GloVe word vector (glove-wiki-gigaword-300) for token k and $q_{i,k}$ is the TF-IDF weight. We compute TF-IDF weights separately for the patent text corpus (fitted separately for each grant year) and the occupation task text corpus. For each token in the document that appears in both the TF-IDF vocabulary and the GloVe model, we multiply the word vector by its TF-IDF weight, sum the weighted vectors, and divide by the number of valid weighted tokens.

Post-processing of similarity scores After computing the cosine similarity $s_{p,j}^r$ between patent p and occupation j ’s routine or non-routine tasks as in Equation (35), we apply two adjustments following [Kogan et al. \(2024\)](#). First, we remove within-year means from the raw cosine similarity scores (separately for routine and non-routine similarities and for each grant year). Second, we impose sparsity by setting values below the 80th percentile to zero, and rescale the remaining values to $[0, 1]$ by dividing by the within-year maximum above the 80th percentile, to obtain the adjusted similarity $\tilde{s}_{p,j}^r$ used in Equation (36).

IA4.2 Similarity between Wikipedia Titles and Occupations’ Task Descriptions

This subsection describes how we construct embeddings for Wikipedia titles used to compute the cosine similarity $s_{w,j}^r$ between Wikipedia title w and occupation j ’s routine or non-routine tasks in Section [VI.A.1](#). The key difference from the patent-based approach in Appendix Section [IA4.1](#) is that we use LLM-extracted technical keywords to represent each Wikipedia title, rather than TF-IDF over the full text. This is because Wikipedia articles contain substantial background information (e.g., time and place of invention) beyond the technology’s functional characteristics, and keywords allow us to focus the embedding on the core technical content.

Keyword extraction from Wikipedia titles For each Wikipedia title and its summary, we use Gemini 2.5 Pro to extract 12 to 25 technical keywords or phrases that capture the core mechanism or principle, key components or materials, main process verbs or operations, and relevant technical domain terms of the technology. The LLM is instructed to prefer domain-specific technical terms over generic words (such as “method,” “system,” or “process”) and to exclude names of people, companies, places, and dates. The prompt is as follows:

You are a technical analyst.

I will provide a Wikipedia title and its summary describing an invention/technology.

Your task is to output a set of technical keywords/phrases that capture:

- core mechanism / principle
- key components / materials
- main process verbs / operations
- relevant technical domain terms

Rules:

- Output ONLY a JSON array of strings. No extra text.
- Include 12 to 25 items.
- Prefer domain-specific technical terms over generic words (avoid: "method", "system", "various", "process", "technology").
- Remove names of people, companies, places, and dates.
- Do NOT invent new information not implied by the input.
- Include both unigrams and short multi-word phrases when appropriate.

Input:

Title: {technology}

Summary:

{summary}

Output:

Wikipedia title embeddings As described in Section VI.A.1, each Wikipedia title w is represented as the simple mean of its keywords’ GloVe word vectors:

$$X_w = \frac{1}{K_w} \sum_{k=1}^{K_w} x_k,$$

where K_w is the number of keyword words for Wikipedia title w and x_k are GloVe word vectors (glove-wiki-gigaword-300). For multi-word keyword phrases, each constituent word is included separately in the average.

Occupation task embeddings The occupation task embeddings \mathcal{X}_j^r ($r \in \{R, NR\}$) are constructed the same way as in the patent-based approach (Appendix Section IA4.1), using TF-IDF-weighted GloVe embeddings of each occupation’s routine and non-routine task descriptions (keeping only verbs and nouns and removing stopwords before calculating the embeddings).

Similarity computation and post-processing We compute cosine similarity $s_{w,j}^r$ between the unweighted Wikipedia title embedding X_w and the TF-IDF weighted occupation task embedding \mathcal{X}_j^r , analogous to Equation (35). We then apply the same post-processing as in the patent-based approach: remove within-year means (separately for routine and non-routine similarities and for each year), impose sparsity at the 80th percentile, and rescale the remaining values to $[0, 1]$ to obtain the adjusted similarity $\tilde{s}_{w,j}^r$ used in Equation (45).

IA5 Measuring the Age of Compustat Firms

For the analyses in Sections III.D, V, and VI, we combine several sources to measure the age of Compustat firms, which we explain in detail below.

IPO date from Compustat We obtain the IPO date (ipodate variable) from Compustat. According to S&P Global Support, Compustat started collecting IPO date in 2002, after which this variable is systematically recorded for companies that have an Initial Public Offering on the NYSE, NYSE America, NYSE ARCA, or Nasdaq. For IPOs before 2002, the date may be missing, or it may be based on other sources or the date when daily pricing began.

Incorporation date from Refinitiv We collect incorporation date from Refinitiv. We merge this information into Compustat based on CUSIP.

LLM query of firm founding year For all unique historical real-time firm names in Compustat (conmh variable), we use Gemini 2.5 Pro to search for the company’s founding year with the following query:

You are an expert in company history.

Return only the *founding year* of this company as a 4-digit number.

If unknown, return "NA".

Company: {company}

Output JSON ONLY:

```

{{
  "company": "{company}",
  "founding_year": "YYYY or NA"
}}

```

Estimated firm age We estimate firm age as the number of years since the earliest year among the available dates above. We set firm age to missing if the earliest of these dates is after the earliest year of the company's record in Compustat, which indicates that our founding year information may not be accurate.

IA6 Appendix for Models

IA6.1 Full Dynamic Model

This section provides a rigorous statement of the dynamic model described in Section III.C. Time is discrete ($t = 0, 1, 2, \dots$) and firm owners discount the future at rate ρ . Business model B has not yet arrived at $t = 0$.

IA6.1.1 State Variables and Timing

At the beginning of each period t , the state of firm i is the length of its rule list $r_{i,t} \in \{0, 1, \dots, R\}$ and whether business model B has arrived in a previous period. The arrival of B is an absorbing, industry-wide event. Denote by τ the (random) period of arrival.

Pre-arrival phase ($t < \tau$). Business model B has not yet arrived. Each period proceeds as follows:

1. *Production under business model A.* All operating units run business model A . Each A -manager solves the static problem (1) and the firm owner chooses $L_{A,i,t}$ to solve (6). Because neither the productivity parameters $\{\psi_{A,j}\}$ nor the manager wage w_A change over time, the per-period solution is time-invariant: $L_{A,i,t} = z_A \xi_i$ and per-period profit from model A is $\pi_A(\xi_i) = \frac{1}{2} z_A^2 \xi_i$.
2. *Arrival draw.* With probability p , business model B arrives ($\tau = t + 1$). With probability $1 - p$, it does not, and the firm proceeds to the incident draw.
3. *Incident draw (conditional on no arrival).* An incident $d_{i,t}$ is drawn uniformly from $\{1, \dots, R\}$:
 - With probability $r_{i,t}/R$, the incident matches an existing rule: it is resolved at no cost, and $r_{i,t+1} = r_{i,t}$.
 - With probability $1 - r_{i,t}/R$, the incident is new: the A -managers pay cost c to resolve it and append a rule, so $r_{i,t+1} = \min(r_{i,t} + 1, R)$.

Post-arrival phase ($t \geq \tau$). Business model B is available. The rule list in business model A stops at $r_i \equiv r_{i,\tau}$. Each period:

1. *Production under business model A.* Unchanged from the pre-arrival phase. No further incidents arise in model A .
2. *Production under business model B.* Each B -manager solves the static problem (9) taking the productivity parameter $\eta_i = \eta - \gamma(r_i) |\Delta|$ as given, where $\gamma(\cdot)$ is derived below. The firm owner chooses $L_{B,i}$ to solve (19). Per-period profit from model B is $\pi_B(r_i, \xi_i) = \frac{1}{2} z_{B,i}^2 \xi_i$, where $z_{B,i} = (\Gamma - \gamma(r_i)) |\Delta| Q_B + L_B / \Xi$ as in (22).

IA6.1.2 Derivation of Organizational Frictions

The key object linking the dynamic rule accumulation to the static production problem is the function $\gamma(r_i)$ in (23). Here we repeat how we microfound this link for completeness.

After B arrives, for each B -operating unit k , an incident arises with probability $z \in (0, 1)$ per period. When an incident arises, A -workers and B -workers agree with probability $1 - |\Delta|$ and disagree with probability $|\Delta|$. Agreement yields a productivity benefit of $z(1 - |\Delta|)\hat{\eta}$ per B -manager, where $\hat{\eta} > 0$. Disagreement triggers a conflict that is escalated to B -managers; each conflict requires $\hat{\gamma}(r_i)$ units of managerial time, with $\hat{\gamma}(\cdot)$ strictly increasing in the rule length.

The total number of conflicts across all $L_{B,i}$ units is $z |\Delta| L_{B,i}$. Shared equally among $L_{B,i}$ managers, each B -manager devotes $z |\Delta| \hat{\gamma}(r_i)$ units of time to conflict resolution. Since each B -manager has one unit of time, the effective productivity of business model B at firm i is:

$$\begin{aligned}
 \eta_i &= 1 + \underbrace{z(1 - |\Delta|)\hat{\eta}}_{\text{benefit of agreement}} - \underbrace{z |\Delta| \hat{\gamma}(r_i)}_{\text{cost of disagreement}} \\
 &= (1 + z\hat{\eta}) - z(\hat{\gamma}(r_i) + \hat{\eta}) |\Delta| \\
 &= \underbrace{\eta}_{\equiv 1 + z\hat{\eta}} - \underbrace{\gamma(r_i)}_{\equiv z(\hat{\gamma}(r_i) + \hat{\eta})} |\Delta|,
 \end{aligned}$$

which is equation (23). The second line follows by regrouping: the constant $1 + z\hat{\eta}$ collects the baseline productivity plus the full agreement benefit, while the term linear in $|\Delta|$ collects both the forgone agreement benefit ($z\hat{\eta} |\Delta|$) and the direct cost of disagreement ($z\hat{\gamma}(r_i) |\Delta|$). Since $\hat{\gamma}(\cdot)$ is strictly increasing, so is $\gamma(\cdot)$: firms with longer rule lists face greater organizational frictions when implementing business model B .

IA6.1.3 Value Functions

Post-arrival value. After B arrives with rule length r_i , the firm collects constant per-period profits from both models in perpetuity:

$$V^{\text{post}}(r_i, \xi_i) = \frac{\pi_A(\xi_i) + \pi_B(r_i, \xi_i)}{\rho}. \quad (\text{IA1})$$

Pre-arrival value. Let $V(r, \xi_i)$ denote the value of firm i at the beginning of a pre-arrival period with r rules. The Bellman equation is:

$$V(r, \xi_i) = \pi_A(\xi_i) + \frac{1}{1 + \rho} \left[p V^{\text{post}}(r, \xi_i) + (1 - p) \left(\frac{r}{R} V(r, \xi_i) + \left(1 - \frac{r}{R} \right) (-c + V(r + 1, \xi_i)) \right) \right], \quad (\text{IA2})$$

with the convention that $V(R + 1, \xi_i) = V(R, \xi_i)$.

IA6.1.4 Key Properties

Three properties of this dynamic model are worth emphasizing.

First, the per-period optimization problems under models A and B are purely static. The occupation composition of each operating unit, the number of operating units, and the equilibrium wages are all determined by the within-period problems in Sections III.A and III.B. The dynamic structure affects only the rule length $r_{i,t}$, which evolves as a stochastic process during the pre-arrival phase and stops upon arrival of B . This separation holds because: (i) rule accumulation is controlled by A -managers, not by the firm owner; (ii) the number of A -units does not affect the incident rate; and (iii) there are no intertemporal adjustment costs.

Second, the arrival of business model B is an industry-wide event: all firms in the industry face the same arrival time τ . Before B arrives, the industry operates solely under model A , and the equilibrium wage $w_{A,t}$ is constant over time—that is, $w_{A,t} = w_A$ for all $t < \tau$ —even though rule lengths $\{r_{i,t}\}$ evolve idiosyncratically across firms. To see this, note that in each pre-arrival period the A -manager's problem (1) and the firm owner's problem (6) are both independent of $r_{i,t}$: rules govern only how incidents are resolved, not the production technology, the occupation composition, or the output Q_A of each operating unit. The firm owner's demand for A -managers is therefore $L_{A,i,t} = (Q_A - w\bar{l} - w_{A,t}) \xi_i$, which depends on ξ_i but not on $r_{i,t}$. Aggregating across firms, market clearing requires $\int_i L_{A,i,t} di = (Q_A - w\bar{l} - w_{A,t}) \Xi = L_A$, which pins down $w_{A,t} = (Q_A - w\bar{l}) - L_A/\Xi \equiv w_A$ independently of t or of the distribution of rules. The per-unit surplus $z_A = L_A/\Xi$ is likewise constant. It follows that the static equilibrium characterization from Section III.A applies unchanged in every pre-arrival period, regardless of how rules are distributed

across firms at that point in time.

Third, the rule length $r_{i,t}$ is (weakly) increasing over time and its expected value increases with firm age. Since $\gamma(\cdot)$ is strictly increasing, older firms—those that have been accumulating rules for longer before B arrives—tend to have higher $\gamma_i = \gamma(r_i)$ and therefore face greater organizational frictions when implementing business model B .

IA6.2 Microfoundation of Conflicts among Divisions

In Section III.C, we provide a microfoundation for why within-firm agency conflicts become more severe when the workstyles associated with the two business models (A and B) are more misaligned (i.e., when $|\Delta|$ is larger). In that setting, the inefficiency originates from disagreements among *workers*, but must ultimately be resolved by managers. We emphasize the worker dimension because a central objective of our framework is to incorporate workstyles into the analysis of firms' abilities to expand their boundaries and implement new businesses.

A natural question is whether firms can avoid the conflicts we model in Section III.C by separating the operation of the new business model B entirely from that of business model A , with no interaction between their workers. In the following, we examine the possibility that firm i establishes a distinct division for business model B , operating separately from the division of business model A , with both divisions reporting to the same headquarters of firm i . In this case, the headquarters of firm i retains majority ownership of division B ; if firm i completely spins off B , then this business would be a separate firm and its growth would not count towards firm i anymore.

Can such separation be successful for overcoming organizational frictions? Examples from practice suggest that companies sometimes entertain this possibility, but there have been very few cases where they have successfully implemented new businesses through internal separation (O'Reilly and Tushman, 2021). In the following, we formalize why it is difficult to have A and B as fully independent entities, as long as they are run by the same headquarters (without selling off B as an independent firm). We consider the agency conflict between divisional managers and the headquarters, which has been studied extensively in the incomplete contracting literature. We show that a natural extension of the incomplete contracting mechanism can deliver the specification $\eta_i = \eta - \gamma_i |\Delta|$ in (18), with the following key features:

$$\frac{\partial \eta_i}{\partial |\Delta|} < 0, \quad \frac{\partial^2 \eta_i}{\partial |\Delta| \partial \gamma_i} < 0. \quad (\text{IA3})$$

That is, a larger workstyle difference $|\Delta|$ reduces the efficiency of business model B , and this negative effect is stronger for older firms (with longer existing rules).

We discuss two types of microfoundations below.

Management aggrievement and shading in Hart and Moore (2008) One type of microfoundations for (IA3) can derive from the notion of “contract as reference point” and the associated “shading” from Hart and Moore (2008). In a nutshell, when A -managers feel “aggrieved” by ex post negative outcomes that fall short from the maximum level specified in the contract, they may engage in shading behavior, such as persuading the headquarters about misbehavior of division B , or delaying information sharing with the headquarters.

More specifically, suppose that the headquarters has a fixed unit (normalized to 1) of resources to allocate between the two divisions. These resources may include operational infrastructure (such as equipment or cloud services) that enhances production efficiency, as well as discretionary budgets that provide employee perks. Let a ($1 - a$) denote the allocation to division A (B); and the A -managers are aggrieved when their realized allocation a falls short of the maximum possible allocation in the contract (say 1). Recall in Section III.E, Result 2, we impose the assumption that $\text{cov}(Q_B | \Delta|, |\Delta|) > 0$. That is to say, industries with larger workstyle changes also tend to have more productive new business models. (Intuitively, firms are undertaking B likely because the new business is sufficiently productive given the organizational cost.) Because a higher Q_B naturally tilts resource allocation to B -division ex post, and $|\Delta|$ is positively correlated with Q_B , the greater the distance between the workstyles of the two business models, the more severe the potential resource shortfall experienced by A -managers. As a result, the A -managers’ grievance, denoted by $g(|\Delta|)$, increases with $|\Delta|$, so that $g'(|\Delta|) > 0$.

Following Hart and Moore (2008), we assume that shading by A -managers reduces the effective efficiency η_i of B -managers, with the following functional form

$$\eta_i = \eta - \gamma_i \cdot g(|\Delta|). \quad (\text{IA4})$$

The interpretation is that manager A can persuade headquarters that division B has engaged in misbehavior, and pressures headquarters investigate division B . Here, η is the baseline efficiency of division B , $g(|\Delta|)$ is the shading activity of division A , and $\gamma_i > 0$ denotes the effectiveness of shading (corresponding to the parameter θ in Hart and Moore (2008)) which can reflect the persuasiveness or influence of division A .

Equation (IA4) delivers our desired property in (IA3). First, because $g'(|\Delta|) > 0$, we have $\frac{\partial \eta_i}{\partial |\Delta|} = -\gamma_i g'(|\Delta|) < 0$. Second, we assume that shading by division A is more powerful—that is, γ_i is larger—when the firm is older, in which case A has accumulated more experiences and more rules (r_i) so their accusations against B are more persuasive and influential; this gives $\frac{\partial^2 \eta_i}{\partial |\Delta| \partial \gamma_i} = -g'(|\Delta|) < 0$. This is because with longer and more detailed rules, it is easier to persuade the headquarters about misbehavior of division B . In sum, the microfoundation along the line of Hart and Moore (2008) generates the organizational friction term η_i analogous to (18) with key properties in (IA3).

More general incomplete contracting mechanism The key properties in (IA3) can also arise from economic forces that are more general than those emphasized by Hart and Moore (2008). For instance, following Scharfstein and Stein (2000) and Rajan, Servaes, and Zingales (2000), when the two divisions share

common budget in the same entity and therefore are not financially independent, the existing division A potentially contributes a significant portion of internal funds and may decide to engage in various forms of sabotage that reduce the efficiency of B , leading to a lower effective productivity parameter η_i . Specifically, without organizational frictions, internal funds accumulated from previous cash flows of A 's operations can be used to increase B 's productivity by η (e.g., providing B with better equipment and facilities). When workstyle difference $|\Delta|$ is larger, division A has more disagreement with how division B operates, and would be more reluctant for headquarters to spend resources on division B . In this case, division A may persuade headquarters to allocate less internal funds for division B , and it can be more powerful when the firm is older in which case A has accumulated more experiences and influence. Accordingly, division B 's efficiency decreases by more—that is, γ_i is larger. Taken together, these forces would imply the presence of an organizational friction term analogous to (18), which inherits the key properties summarized in (IA3).

IA6.3 Proofs for Connection to Empirical Analyses

IA6.3.1 Proof of Result 1

Proof. We work within a single industry and suppress the industry subscript n where no confusion arises. From the model in Sections III.A and III.B, pre-period employment at firm i is $E_{\text{pre},i} = z_A \xi_i \bar{l}$, where $z_A = L_A/\Xi$. Post-period employment is $E_{\text{post},i} = (z_A + z_{B,i}) \xi_i \bar{l}$, where $z_{B,i} = (\Gamma - \gamma_i) |\Delta| Q_B + L_B/\Xi$. The growth rate of firm i is therefore:

$$g_i = \frac{E_{\text{post},i} - E_{\text{pre},i}}{E_{\text{pre},i}} = \frac{z_{B,i}}{z_A} = \frac{(\Gamma - \gamma_i) |\Delta| Q_B + L_B/\Xi}{L_A/\Xi}.$$

Note that ξ_i cancels between numerator and denominator, so growth is independent of firm size. Rearranging:

$$g_i = \frac{L_B}{L_A} + \frac{\Xi}{L_A} Q_B |\Delta| \Gamma - \frac{\Xi}{L_A} Q_B |\Delta| \gamma_i = \mu - \lambda Q_B |\Delta| \gamma_i,$$

where $\lambda \equiv \Xi/L_A$ and $\mu \equiv L_B/L_A + \lambda Q_B |\Delta| \Gamma$.

For the workstyle change, the pre-period industry average workstyle is:

$$\bar{w}S_{\text{pre}} = \sum_j \frac{\psi_{A,j}}{\sum_{j'} \psi_{A,j'}} wS_j,$$

since all firms have the same occupation composition under business model A . In the post-period:

$$\bar{w}S_{\text{post}} = \bar{w}S_{\text{pre}} + \frac{L_B}{L_A + L_B} \sum_j \left(\frac{\psi_{B,j}}{\sum_{j'} \psi_{B,j'}} - \frac{\psi_{A,j}}{\sum_{j'} \psi_{A,j'}} \right) wS_j,$$

where the weight $L_B/(L_A + L_B)$ reflects the share of total employment under business model B . Thus

$$|d\bar{w}s| = \frac{L_B}{L_A + L_B} |\Delta|. \quad \square$$

IA6.3.2 Proof of Result 2

We restate the result here for convenience. To make notation lighter, write $w_n \equiv |d\bar{w}s|_n$ for the absolute change in industry-level workstyle. Recall the equilibrium conditions of the model:

$$g_{i,n} = \mu_n - \lambda Q_{B,n} |\Delta|_n \gamma_i, \quad (\text{IA5})$$

$$\lambda = \frac{\Xi}{L_A}, \quad (\text{IA6})$$

$$\mu_n = \frac{L_B}{L_A} + \lambda Q_{B,n} |\Delta|_n \Gamma, \quad (\text{IA7})$$

$$w_n = \Theta |\Delta|_n, \quad (\text{IA8})$$

$$\Theta = \frac{L_B}{L_A + L_B}. \quad (\text{IA9})$$

We maintain the following assumption, restated from the main text for convenience.

Assumption IA1 (Assumption 1, restated). *The joint distribution of (age_i, γ_i) is identical across industries n , and is independent of all industry-level variables (including $Q_{B,n}$, $|\Delta|_n$, and $|d\bar{w}s|_n$).*

Result IA3 (Result 2, restated). *Under Assumption IA1, the population OLS coefficient on $w_n \times \text{age}_i$ in regression (29) is*

$$\hat{\zeta} = -\Phi \times \frac{\text{cov}(\gamma_i, \text{age}_i)}{\text{var}(\text{age}_i)}, \quad (\text{IA10})$$

where

$$\Phi \equiv \frac{\lambda}{\Theta} \frac{\text{cov}(Q_{B,n} |\Delta|_n, |\Delta|_n)}{\text{var}(|\Delta|_n)} \quad (\text{IA11})$$

is an industry-level scalar expressed in terms of the equilibrium objects in (IA5)–(IA9). Sufficient conditions for $\Phi > 0$ are given in Section IA6.3.3.

Proof. We apply the Frisch–Waugh–Lovell (FWL) theorem. Consider the regression

$$g_{i,n} = \nu + \kappa w_n + \rho \text{age}_i + \zeta (w_n \times \text{age}_i) + \varepsilon_{i,n}. \quad (\text{IA12})$$

The FWL theorem states that the coefficient ζ on the interaction term $w_n \times \text{age}_i$ in (IA12) can be obtained equivalently by: (i) regressing $w_n \times \text{age}_i$ on the other regressors $(1, w_n, \text{age}_i)$ and collecting the residual; (ii) regressing $g_{i,n}$ on the same regressors $(1, w_n, \text{age}_i)$ and collecting the residual; and (iii) running a univariate regression of the residual from (ii) on the residual from (i). We now carry out each of these three regressions explicitly.

Regression (i): projecting $w_n \times \text{age}_i$ onto $(1, w_n, \text{age}_i)$. Consider the population regression

$$w_n \times \text{age}_i = \alpha_1 + \beta_2 w_n + \beta_1 \text{age}_i + r_{i,n}^{(1)}, \quad (\text{IA13})$$

where $r_{i,n}^{(1)}$ is the residual. The population OLS coefficients $(\alpha_1, \beta_1, \beta_2)$ satisfy the normal equations. Since Assumption IA1 implies that firm-level and industry-level variables are independent, $\text{cov}(\text{age}_i, w_n) = 0$, and the normal equations for (β_2, β_1) reduce to the diagonal system

$$\begin{pmatrix} \text{var}(w_n) & 0 \\ 0 & \text{var}(\text{age}_i) \end{pmatrix} \begin{pmatrix} \beta_2 \\ \beta_1 \end{pmatrix} = \begin{pmatrix} \text{cov}(w_n \times \text{age}_i, w_n) \\ \text{cov}(w_n \times \text{age}_i, \text{age}_i) \end{pmatrix}.$$

Using firm–industry independence to evaluate the right-hand side:

$$\text{cov}(w_n \times \text{age}_i, w_n) = \mathbb{E}(\text{age}_i) \text{var}(w_n),$$

$$\text{cov}(w_n \times \text{age}_i, \text{age}_i) = \mathbb{E}(w_n) \text{var}(\text{age}_i),$$

so $\beta_2 = \mathbb{E}(\text{age}_i)$ and $\beta_1 = \mathbb{E}(w_n)$. The intercept is $\alpha_1 = \mathbb{E}(w_n \times \text{age}_i) - \beta_2 \mathbb{E}(w_n) - \beta_1 \mathbb{E}(\text{age}_i) = -\mathbb{E}(w_n) \mathbb{E}(\text{age}_i)$, where we used $\mathbb{E}(w_n \times \text{age}_i) = \mathbb{E}(w_n) \mathbb{E}(\text{age}_i)$ (again by independence). Substituting back into (IA13), the residual is:

$$\begin{aligned} r_{i,n}^{(1)} &= w_n \times \text{age}_i + \mathbb{E}(w_n) \mathbb{E}(\text{age}_i) - \mathbb{E}(\text{age}_i) w_n - \mathbb{E}(w_n) \text{age}_i \\ &= (w_n - \mathbb{E}(w_n)) (\text{age}_i - \mathbb{E}(\text{age}_i)) \equiv v_n a_i, \end{aligned} \quad (\text{IA14})$$

where we define $v_n \equiv w_n - \mathbb{E}(w_n)$ and $a_i \equiv \text{age}_i - \mathbb{E}(\text{age}_i)$.

Regression (ii): projecting $g_{i,n}$ onto $(1, w_n, \text{age}_i)$. Consider the population regression

$$g_{i,n} = \alpha_2 + \kappa_2 w_n + \rho_2 \text{age}_i + r_{i,n}^{(2)}, \quad (\text{IA15})$$

where $r_{i,n}^{(2)}$ is the residual. We do not need to solve for $(\alpha_2, \rho_2, \kappa_2)$ explicitly. What matters is a property of the residual $r_{i,n}^{(2)}$: by the definition of population OLS, $r_{i,n}^{(2)}$ is uncorrelated with each of the regressors $(1, w_n, \text{age}_i)$.

Regression (iii): univariate regression of $r_{i,n}^{(2)}$ on $r_{i,n}^{(1)}$. By the FWL theorem, the coefficient ζ in (IA12) equals the coefficient from the univariate regression of $r_{i,n}^{(2)}$ on $r_{i,n}^{(1)} = v_n a_i$:

$$\zeta = \frac{\text{cov}(r_{i,n}^{(2)}, v_n a_i)}{\text{var}(v_n a_i)}. \quad (\text{IA16})$$

We now show that $\text{cov}(r_{i,n}^{(2)}, v_n a_i) = \text{cov}(g_{i,n}, v_n a_i)$, so that $r_{i,n}^{(2)}$ can be replaced by $g_{i,n}$ in the numerator.

From (IA15), $r_{i,n}^{(2)} = g_{i,n} - \alpha_2 - \rho_2 \text{age}_i - \kappa_2 w_n$, so

$$\text{cov}(r_{i,n}^{(2)}, v_n a_i) = \text{cov}(g_{i,n}, v_n a_i) - \rho_2 \text{cov}(\text{age}_i, v_n a_i) - \kappa_2 \text{cov}(w_n, v_n a_i).$$

Both $\text{cov}(\text{age}_i, v_n a_i)$ and $\text{cov}(w_n, v_n a_i)$ are zero, because $v_n a_i = r_{i,n}^{(1)}$ is the residual from regression (i) and is therefore orthogonal to $(1, w_n, \text{age}_i)$. This gives

$$\zeta = \frac{\text{cov}(g_{i,n}, v_n a_i)}{\text{var}(v_n a_i)}. \quad (\text{IA17})$$

It remains to evaluate the numerator and denominator of (IA17) using the equilibrium conditions and Assumption IA1. For the denominator, since v_n depends only on industry n and a_i depends only on firm i , firm–industry independence gives

$$\text{var}(v_n a_i) = \text{var}(w_n) \text{var}(\text{age}_i).$$

For the numerator, substituting (IA5):

$$\text{cov}(g_{i,n}, v_n a_i) = \text{cov}(\mu_n - \lambda Q_{B,n} |\Delta|_n \gamma_i, v_n a_i).$$

The first term μ_n depends only on industry n , while a_i has mean zero and is independent of v_n , so $\text{cov}(\mu_n, v_n a_i) = \mathbb{E}(\mu_n v_n) \mathbb{E}(a_i) = 0$. For the second term, we use firm–industry independence to factor the expectation:

$$\begin{aligned} \text{cov}(Q_{B,n} |\Delta|_n \gamma_i, v_n a_i) &= \mathbb{E}(Q_{B,n} |\Delta|_n \gamma_i v_n a_i) - \mathbb{E}(Q_{B,n} |\Delta|_n \gamma_i) \mathbb{E}(v_n a_i) \\ &= \mathbb{E}(Q_{B,n} |\Delta|_n v_n) \mathbb{E}(\gamma_i a_i) - \mathbb{E}(Q_{B,n} |\Delta|_n) \mathbb{E}(\gamma_i) \mathbb{E}(v_n) \mathbb{E}(a_i) \\ &= \mathbb{E}(Q_{B,n} |\Delta|_n v_n) \text{cov}(\gamma_i, \text{age}_i), \end{aligned}$$

where the last line uses $\mathbb{E}(a_i) = 0$ and $\mathbb{E}(\gamma_i a_i) = \text{cov}(\gamma_i, \text{age}_i)$. Next, from (IA8), $v_n = w_n - \mathbb{E}(w_n) = \Theta(|\Delta|_n - \mathbb{E}(|\Delta|_n))$, so

$$\mathbb{E}(Q_{B,n} |\Delta|_n v_n) = \Theta \text{cov}(Q_{B,n} |\Delta|_n, |\Delta|_n).$$

Combining, the numerator of (IA17) is:

$$\text{cov}(g_{i,n}, v_n a_i) = -\lambda \Theta \text{cov}(Q_{B,n} |\Delta|_n, |\Delta|_n) \text{cov}(\gamma_i, \text{age}_i). \quad (\text{IA18})$$

Substituting (IA18) and $\text{var}(w_n) = \Theta^2 \text{var}(|\Delta|_n)$ into (IA17):

$$\zeta = \frac{-\lambda \Theta \text{cov}(Q_{B,n} | \Delta|_n, |\Delta|_n) \text{cov}(\gamma_i, \text{age}_i)}{\text{var}(\text{age}_i) \Theta^2 \text{var}(|\Delta|_n)} = -\Phi \times \frac{\text{cov}(\gamma_i, \text{age}_i)}{\text{var}(\text{age}_i)},$$

where Φ is as defined in (IA11). □

IA6.3.3 Sufficient Conditions for $\Phi > 0$

Assumption IA2. *The occupation productivity vectors $\{\psi_{A,j,n}\}$ and $\{\psi_{B,j,n}\}$ across industries can be written as:*

$$\begin{aligned} \psi_{A,j,n} &= s_{A,j,n} \psi_n, \\ \psi_{B,j,n} &= s_{B,j,n} (1 + \nu_n) \psi_n, \end{aligned} \tag{IA19}$$

where the occupation share vectors $\{s_{A,j,n}\}$ and $\{s_{B,j,n}\}$ are independently distributed from the scalars ψ_n and ν_n , and satisfy $\sum_j s_{A,j,n} = \sum_j s_{B,j,n} = 1$.

Economically, Assumption IA2 says that the new business model B scales the occupation productivity aggregate $\sum_j \psi_{B,j,n}$ relative to $\sum_j \psi_{A,j,n}$ by an industry-specific factor $(1 + \nu_n)$, implying an output ratio $Q_{B,n}/Q_{A,n} = (1 + \nu_n)^{1/(\sigma-1)}$. Model B also alters the composition of occupations required relative to A , as captured by the occupation share vectors $\{s_{A,j,n}\}$ and $\{s_{B,j,n}\}$. Crucially, the *scale* of the productivity improvement (governed by ψ_n and ν_n) is independent of the *direction* of the occupational recomposition (governed by the share vectors). Under this decomposition, the output per unit in the new business model, $Q_{B,n}$, depends only on (ψ_n, ν_n) , while the workstyle distance $|\Delta|_n$ depends only on the share vectors. This independence ensures that $\Phi > 0$.

Result IA4. *Under Assumption IA2, $\Phi > 0$. Specifically:*

$$\Phi = \frac{\lambda}{\Theta} \mathbb{E}(Q_{B,n}) = \frac{\Theta}{L_A} \left(1 + \frac{L_A}{L_B}\right) \times \mathbb{E}\left[\left((1 + \nu_n) \psi_n\right)^{\frac{1}{\sigma-1}}\right] \bar{l} > 0. \tag{IA20}$$

Proof. Under Assumption IA2:

$$\begin{aligned} Q_{B,n} &= (1 + \nu_n)^{\frac{1}{\sigma-1}} \psi_n^{\frac{1}{\sigma-1}} \bar{l}, \\ |\Delta|_n &= \left| \sum_j (s_{B,j,n} - s_{A,j,n}) w s_j \right|. \end{aligned}$$

Since $Q_{B,n}$ depends only on (ψ_n, ν_n) and $|\Delta|_n$ depends only on $\{s_{A,j,n}, s_{B,j,n}\}$, the two are independent. Therefore:

$$\text{cov}(Q_{B,n} | \Delta|_n, |\Delta|_n) = \mathbb{E}(Q_{B,n}) \text{var}(|\Delta|_n).$$

Substituting into the definition of Φ in (IA11):

$$\Phi = \frac{\lambda}{\Theta} \frac{\mathbb{E}(Q_{B,n}) \text{var}(|\Delta|_n)}{\text{var}(|\Delta|_n)} = \frac{\lambda}{\Theta} \mathbb{E}(Q_{B,n}) > 0,$$

since $\lambda, \Theta > 0$ and $Q_{B,n} > 0$. □

More general conditions for $\Phi > 0$ Assumption IA2 is sufficient but not necessary. More generally, from the definition in (IA11), $\Phi > 0$ if and only if

$$\text{cov}(Q_{B,n} |\Delta|_n, |\Delta|_n) > 0, \quad (\text{IA21})$$

since $\lambda/\Theta > 0$. The left-hand side is the population covariance, across industries, between the “effective disruption” $Q_{B,n} |\Delta|_n$ and the workstyle distance $|\Delta|_n$. Condition (IA21) therefore holds whenever industries with larger workstyle changes also tend to have higher effective disruption, that is, workstyle-disruptive technologies tend to also be productive ones. This condition is empirically testable: it can be assessed by regressing a proxy for $Q_{B,n} |\Delta|_n$ on $|\Delta|_n$ in the cross-section of industries.

To see that (IA21) holds under weak conditions even without full independence, define $|\widetilde{\Delta}|_n \equiv |\Delta|_n - \mathbb{E}(|\Delta|_n)$ and decompose the covariance as

$$\text{cov}(Q_{B,n} |\Delta|_n, |\Delta|_n) = \underbrace{\mathbb{E}(Q_{B,n}) \text{var}(|\Delta|_n)}_{>0} + \underbrace{\text{cov}(Q_{B,n}, |\Delta|_n |\widetilde{\Delta}|_n)}_{\geq 0}. \quad (\text{IA22})$$

The first term is strictly positive. The correction term vanishes under Assumption IA2 (independence), and is small when $Q_{B,n}$ and $|\Delta|_n$ are weakly dependent. The condition $\Phi > 0$ therefore fails only if the correction is sufficiently negative to outweigh the positive baseline $\mathbb{E}(Q_{B,n}) \text{var}(|\Delta|_n)$ —meaning the most productive new business models (high $Q_{B,n}$) happen to require the *least* workstyle change (low $|\Delta|_n$). This is the economically implausible case in which adopting business model B yields high output but requires little organizational adjustment, so that there is no friction-based disadvantage for older firms.

IA6.3.4 Regression with Controls

We also generalize Result IA3 to a regression that includes an additional industry-level control Z_n and an additional firm-level control x_i , along with all pairwise interactions of industry-level and firm-level variables. Adding controls does not change the interpretation of the key interaction coefficient ζ_{wa} , but replaces the bivariate coefficients in Result IA3 with their multivariate (conditional) counterparts.

Notation We maintain the notation from the proof of Result IA3: $w_n \equiv |d\overline{w}_s|_n$ denotes the absolute change in industry-level workstyle, and the equilibrium conditions (IA5)–(IA8) carry over unchanged. Let Z_n denote an additional industry-level control and x_i an additional firm-level control. Define the

demeaned variables:

$$\begin{aligned} v_n &\equiv w_n - \mathbb{E}(w_n), & u_n &\equiv Z_n - \mathbb{E}(Z_n), \\ a_i &\equiv \text{age}_i - \mathbb{E}(\text{age}_i), & s_i &\equiv x_i - \mathbb{E}(x_i). \end{aligned}$$

Let Σ_I denote the variance–covariance matrix of the industry-level variables (w_n, Z_n) , and Σ_F the variance–covariance matrix of the firm-level variables (age_i, x_i) :

$$\Sigma_I = \begin{pmatrix} \text{var}(w_n) & \text{cov}(w_n, Z_n) \\ \text{cov}(w_n, Z_n) & \text{var}(Z_n) \end{pmatrix}, \quad \Sigma_F = \begin{pmatrix} \text{var}(\text{age}_i) & \text{cov}(\text{age}_i, x_i) \\ \text{cov}(\text{age}_i, x_i) & \text{var}(x_i) \end{pmatrix}. \quad (\text{IA23})$$

Consider the population multivariate regression of $Q_{B,n} \mid \Delta|_n$ on (w_n, Z_n) :

$$Q_{B,n} \mid \Delta|_n = \alpha_\Phi + \Phi_{w|Z} w_n + \Phi_{Z|w} Z_n + u_n^{(\Phi)}. \quad (\text{IA24})$$

By definition of population OLS, the coefficient vector $(\Phi_{w|Z}, \Phi_{Z|w})'$ satisfies

$$\Sigma_I \begin{pmatrix} \Phi_{w|Z} \\ \Phi_{Z|w} \end{pmatrix} = \begin{pmatrix} \text{cov}(Q_{B,n} \mid \Delta|_n, w_n) \\ \text{cov}(Q_{B,n} \mid \Delta|_n, Z_n) \end{pmatrix}. \quad (\text{IA25})$$

Similarly, consider the population multivariate regression of the organizational friction γ_i on (age_i, x_i) :

$$\gamma_i = \alpha_\gamma + \beta_{a|x} \text{age}_i + \beta_{x|a} x_i + u_i^{(\gamma)}. \quad (\text{IA26})$$

The coefficient vector $(\beta_{a|x}, \beta_{x|a})'$ satisfies

$$\Sigma_F \begin{pmatrix} \beta_{a|x} \\ \beta_{x|a} \end{pmatrix} = \begin{pmatrix} \text{cov}(\gamma_i, \text{age}_i) \\ \text{cov}(\gamma_i, x_i) \end{pmatrix}. \quad (\text{IA27})$$

Assumption IA3. *The joint distribution of $(\text{age}_i, \gamma_i, x_i)$ is identical across industries n , and is independent of all industry-level variables (including $Q_{B,n} \mid \Delta|_n$, w_n , and Z_n). No restriction is imposed on $\text{cov}(w_n, Z_n)$ or $\text{cov}(\text{age}_i, x_i)$.*

Result IA5. *Under Assumption IA3, the population OLS coefficient on $w_n \times \text{age}_i$ in the extended regression*

$$\begin{aligned} g_{i,n} &= \nu + \kappa_w w_n + \kappa_Z Z_n + \rho_a \text{age}_i + \rho_x x_i \\ &\quad + \zeta_{wa} (w_n \times \text{age}_i) + \zeta_{wx} (w_n \times x_i) + \zeta_{Za} (Z_n \times \text{age}_i) + \zeta_{Zx} (Z_n \times x_i) + \varepsilon_{i,n} \end{aligned} \quad (\text{IA28})$$

is

$$\zeta_{wa} = -\lambda \Phi_{w|Z} \times \beta_{a|x}, \quad (\text{IA29})$$

where $\Phi_{w|Z}$ is the coefficient on w_n in the population regression of $Q_{B,n} \mid \Delta|_n$ on (w_n, Z_n) in (IA24), and $\beta_{a|x}$ is

the coefficient on age_i in the population regression of γ_i on (age_i, x_i) in (IA26). More generally, the full 4×1 coefficient vector on the interactions satisfies

$$\begin{pmatrix} \zeta_{wa} \\ \zeta_{wx} \\ \zeta_{Za} \\ \zeta_{Zx} \end{pmatrix} = -\lambda \begin{pmatrix} \Phi_{w|Z} \\ \Phi_{Z|w} \end{pmatrix} \otimes \begin{pmatrix} \beta_{a|x} \\ \beta_{x|a} \end{pmatrix}, \quad (\text{IA30})$$

where \otimes denotes the Kronecker product.

Relative to Result IA3, the key changes are: (i) the industry-level scalar Φ is replaced by $\lambda \Phi_{w|Z}$, the conditional regression coefficient of $Q_{B,n} |\Delta|_n$ on w_n given Z_n ; and (ii) the ratio $\text{cov}(\gamma_i, \text{age}_i) / \text{var}(\text{age}_i)$ is replaced by $\beta_{a|x}$, the conditional regression coefficient of γ_i on age_i given x_i .

Proof. We apply the Frisch–Waugh–Lovell (FWL) theorem in three steps: (i) project each of the four interaction terms onto the level regressors $(1, w_n, Z_n, \text{age}_i, x_i)$ and collect residuals; (ii) project $g_{i,n}$ onto the same level regressors and collect the residual; (iii) run a multivariate regression of the residual from (ii) on the four residuals from (i).

Regression (i): projecting interaction terms onto level regressors.

The variance matrix of the demeaned level regressors (v_n, u_n, a_i, s_i) is block-diagonal by firm–industry independence:

$$\text{var} \begin{pmatrix} v_n \\ u_n \\ a_i \\ s_i \end{pmatrix} = \begin{pmatrix} \Sigma_I & 0 \\ 0 & \Sigma_F \end{pmatrix}.$$

For each interaction (industry variable) \times (firm variable), the right-hand side of the normal equations factors by firm–industry independence. Take $w_n \times \text{age}_i$ as an example. Its covariances with the level regressors are:

$$\text{cov}(w_n \times \text{age}_i, w_n) = \mathbb{E}(\text{age}_i) \text{var}(w_n), \quad \text{cov}(w_n \times \text{age}_i, Z_n) = \mathbb{E}(\text{age}_i) \text{cov}(w_n, Z_n),$$

$$\text{cov}(w_n \times \text{age}_i, \text{age}_i) = \mathbb{E}(w_n) \text{var}(\text{age}_i), \quad \text{cov}(w_n \times \text{age}_i, x_i) = \mathbb{E}(w_n) \text{cov}(\text{age}_i, x_i).$$

The industry block of the RHS is $\mathbb{E}(\text{age}_i)$ times the first column of Σ_I , giving coefficients $(\mathbb{E}(\text{age}_i), 0)$ on (w_n, Z_n) . The firm block is $\mathbb{E}(w_n)$ times the first column of Σ_F , giving coefficients $(\mathbb{E}(w_n), 0)$ on (age_i, x_i) . The residual is therefore:

$$r_{i,n}^{(w,a)} = v_n a_i. \quad (\text{IA31})$$

The same block-diagonal argument applies to each of the remaining three interactions:

$$r_{i,n}^{(w,x)} = v_n s_i, \quad r_{i,n}^{(Z,a)} = u_n a_i, \quad r_{i,n}^{(Z,x)} = u_n s_i. \quad (\text{IA32})$$

Regression (ii): projecting $g_{i,n}$ onto level regressors.

The residual $r_{i,n}^{(g)}$ from the population regression of $g_{i,n}$ on $(1, w_n, Z_n, \text{age}_i, x_i)$ is, by construction, uncorrelated with each of the level regressors.

Regression (iii): multivariate regression of $r_{i,n}^{(g)}$ on interaction residuals.

By the FWL theorem, the coefficient vector $(\zeta_{wa}, \zeta_{wx}, \zeta_{Za}, \zeta_{Zx})'$ equals the coefficient from regressing $r_{i,n}^{(g)}$ on $(v_n a_i, v_n s_i, u_n a_i, u_n s_i)$. As in the proof of Result IA3, the interaction residuals are orthogonal to the level regressors, so $r_{i,n}^{(g)}$ can be replaced by $g_{i,n}$ on the right-hand side. The normal equations are:

$$M \begin{pmatrix} \zeta_{wa} \\ \zeta_{wx} \\ \zeta_{Za} \\ \zeta_{Zx} \end{pmatrix} = \begin{pmatrix} \text{COV}(g_{i,n}, v_n a_i) \\ \text{COV}(g_{i,n}, v_n s_i) \\ \text{COV}(g_{i,n}, u_n a_i) \\ \text{COV}(g_{i,n}, u_n s_i) \end{pmatrix}, \quad (\text{IA33})$$

where M is the 4×4 variance-covariance matrix of $(v_n a_i, v_n s_i, u_n a_i, u_n s_i)$.

Evaluating the left-hand side By firm-industry independence, each entry of M factors into an industry-level moment times a firm-level moment. Specifically, for industry-level indices $r, t \in \{v, u\}$ and firm-level indices $p, q \in \{a, s\}$:

$$\text{COV}(r_n p_i, t_n q_i) = \text{COV}(r_n, t_n) \text{COV}(p_i, q_i).$$

This gives $M = \Sigma_I \otimes \Sigma_F$, where the rows and columns are ordered as $(v \cdot a, v \cdot s, u \cdot a, u \cdot s)$ with the convention (r, p) corresponding to the (r, p) -th element of the Kronecker product.

Evaluating the right-hand side Substituting $g_{i,n} = \mu_n - \lambda Q_{B,n} |\Delta|_n \gamma_i$ from (IA5), and using $\text{COV}(\mu_n, r_n p_i) = 0$ (since $\mathbb{E}(p_i) = 0$ and μ_n is industry-level), we have for each industry index $r \in \{v, u\}$ and firm index $p \in \{a, s\}$:

$$\text{COV}(g_{i,n}, r_n p_i) = -\lambda \mathbb{E}(Q_{B,n} |\Delta|_n r_n) \text{COV}(\gamma_i, p_i),$$

where firm-industry independence is used to factor the expectation. Stacking these into a vector:

$$\begin{pmatrix} \text{COV}(g_{i,n}, v_n a_i) \\ \text{COV}(g_{i,n}, v_n s_i) \\ \text{COV}(g_{i,n}, u_n a_i) \\ \text{COV}(g_{i,n}, u_n s_i) \end{pmatrix} = -\lambda \begin{pmatrix} \text{COV}(Q_{B,n} |\Delta|_n, w_n) \\ \text{COV}(Q_{B,n} |\Delta|_n, Z_n) \end{pmatrix} \otimes \begin{pmatrix} \text{COV}(\gamma_i, \text{age}_i) \\ \text{COV}(\gamma_i, x_i) \end{pmatrix}, \quad (\text{IA34})$$

where we used $\mathbb{E}(Q_{B,n} | \Delta|_n v_n) = \text{cov}(Q_{B,n} | \Delta|_n, w_n)$ and $\mathbb{E}(Q_{B,n} | \Delta|_n u_n) = \text{cov}(Q_{B,n} | \Delta|_n, Z_n)$.

Solving the system Substituting into (IA33) and using $M^{-1} = \Sigma_I^{-1} \otimes \Sigma_F^{-1}$:

$$\begin{pmatrix} \zeta_{wa} \\ \zeta_{wx} \\ \zeta_{Za} \\ \zeta_{Zx} \end{pmatrix} = -\lambda \left(\Sigma_I^{-1} \begin{pmatrix} \text{cov}(Q_{B,n} | \Delta|_n, w_n) \\ \text{cov}(Q_{B,n} | \Delta|_n, Z_n) \end{pmatrix} \right) \otimes \left(\Sigma_F^{-1} \begin{pmatrix} \text{cov}(\gamma_i, \text{age}_i) \\ \text{cov}(\gamma_i, x_i) \end{pmatrix} \right). \quad (\text{IA35})$$

By (IA25), $\Sigma_I^{-1}(\text{cov}(Q_{B,n} | \Delta|_n, w_n), \text{cov}(Q_{B,n} | \Delta|_n, Z_n))' = (\Phi_{w|Z}, \Phi_{Z|w})'$. By (IA27):

$$\Sigma_F^{-1}(\text{cov}(\gamma_i, \text{age}_i), \text{cov}(\gamma_i, x_i))' = (\beta_{a|x}, \beta_{x|a})'.$$

Therefore:

$$\begin{pmatrix} \zeta_{wa} \\ \zeta_{wx} \\ \zeta_{Za} \\ \zeta_{Zx} \end{pmatrix} = -\lambda \begin{pmatrix} \Phi_{w|Z} \\ \Phi_{Z|w} \end{pmatrix} \otimes \begin{pmatrix} \beta_{a|x} \\ \beta_{x|a} \end{pmatrix},$$

which gives (IA30). In particular, $\zeta_{wa} = -\lambda \Phi_{w|Z} \beta_{a|x}$, which is (IA29). \square

IA6.3.5 Firm-Level Shares of Business Model B Output

The following result characterizes firm i 's share of total business model B output in the industry, and how it depends on organizational frictions γ_i and firm scale ξ_i .

Result IA6. *In each industry n , the share of total business model B output produced by firm i is*

$$\sigma_{B,i} = \frac{\xi_i}{\Xi} \omega(\gamma_i), \quad \omega(\gamma_i) \equiv \frac{\eta_i z_{B,i}}{\bar{D}}, \quad (\text{IA36})$$

where $\eta_i = \eta - \gamma_i |\Delta|$, $z_{B,i} = (\Gamma - \gamma_i) |\Delta| Q_B + L_B / \Xi$, and $\bar{D} \equiv \int_j \eta_j z_{B,j} (\xi_j / \Xi) dj$ is an industry-level constant.

The share $\sigma_{B,i}$ has the following properties:

- (i) $\sigma_{B,i}$ is strictly decreasing in γ_i : firms with greater organizational frictions produce a smaller share of business model B output.
- (ii) $\sigma_{B,i}$ is strictly increasing in ξ_i : larger firms (in terms of adjustment cost capacity) produce a larger share.
- (iii) The ratio of firm i 's share of B output to its share of A output is

$$\frac{\sigma_{B,i}}{\sigma_{A,i}} = \omega(\gamma_i), \quad (\text{IA37})$$

which depends only on γ_i , not on ξ_i . This ratio is strictly decreasing in γ_i .

Proof. Firm i 's total output under business model B is $Y_{B,i} = \eta_i Q_B L_{B,i}$, since each of its $L_{B,i}$ operating units produces $\eta_i Q_B$. Total industry B output is $Y_B = Q_B \int_j \eta_j L_{B,j} dj$. Using $L_{B,i} = z_{B,i} \xi_i$:

$$\sigma_{B,i} = \frac{\eta_i z_{B,i} \xi_i}{\int_j \eta_j z_{B,j} \xi_j dj} = \frac{\xi_i}{\Xi} \cdot \frac{\eta_i z_{B,i}}{\int_j \eta_j z_{B,j} (\xi_j/\Xi) dj} = \frac{\xi_i}{\Xi} \omega(\gamma_i),$$

which is (IA36). Under business model A , all firms face the same per-unit surplus z_A , so $L_{A,i} = z_A \xi_i$ and $\sigma_{A,i} = \xi_i/\Xi$. Hence $\sigma_{B,i}/\sigma_{A,i} = \omega(\gamma_i)$, which is (iii).

For (i), differentiate $\eta_i z_{B,i}$ with respect to γ_i :

$$\frac{\partial(\eta_i z_{B,i})}{\partial \gamma_i} = -|\Delta| z_{B,i} + \eta_i (-|\Delta| Q_B) = -|\Delta| (z_{B,i} + \eta_i Q_B) < 0,$$

since $z_{B,i} > 0$ and $\eta_i > 0$. For (ii), note that $\sigma_{B,i}$ is linear in ξ_i with a positive coefficient $\omega(\gamma_i)/\Xi > 0$. \square

The factorization $\sigma_{B,i} = (\xi_i/\Xi) \omega(\gamma_i)$ separates the firm's share of business model B output into two components. The first, ξ_i/Ξ , is the firm's pre-existing share of business model A and reflects its scale advantage. The second, $\omega(\gamma_i)$, captures how organizational frictions shift the firm's relative position when transitioning to B . The strictly decreasing relationship between $\omega(\gamma_i)$ and γ_i reflects two reinforcing channels: higher organizational frictions reduce both the per-unit efficiency of B -managers (through η_i) and the number of B -units the firm operates (through $z_{B,i}$).

Since the ratio $\sigma_{B,i}/\sigma_{A,i}$ is independent of ξ_i , firm scale per se does not determine whether a firm implements business model B more than A . It is organizational frictions alone that govern the reallocation of activity from A to B relative to the firm's pre-existing size. In particular, if organizational frictions γ_i are positively correlated with firm age (as in the microfoundation in Section III.C), young firms will capture a disproportionately large share of business model B output relative to their pre-existing scale, even if they are small in absolute terms.