NBER WORKING PAPER SERIES

AI ADOPTION IN AMERICA: WHO, WHAT, AND WHERE

Kristina McElheran ⓡ
J. Frank Li ⓡ
Erik Brynjolfsson ⓡ
Zachary Kroff ⓡ
Emin Dinlersoz ⓡ
Lucia S. Foster ⓡ
Nikolas Zolas ⓡ

Working Paper 31788
http://www.nber.org/papers/w31788

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 2023

Any opinions and conclusions expressed herein are those of the authors and do not represent the views of the U.S. Census Bureau or the National Bureau of Economic Research. This study relies on confidential micro-data available to approved projects through the U.S. Federal Statistical Research Data Centers. All results have been reviewed to ensure that no confidential information is disclosed. The Census Bureau has reviewed this data product to ensure appropriate access, use, and disclosure avoidance protection of the confidential source data (Project No. P-7508509, Disclosure Review Board (DRB) approval number: CBDRB-FY20-095, CBDRB-FY20-331, CBDRB-FY21-041, CBDRB-FY22-074. CBDRB-FY22-246. CBDRB-FY23-0309). We thank the Editor, two anonymous referees, Tim Bresnahan, John Eltinge, Nathan Goldschlag, Rob Seamans, and participants in the Columbia/Wharton Management, Analytics, and Data Conference and Georgetown's McDonough School of Business strategy seminar for excellent comments and feedback. Financial support from the Stanford Digital Economy Lab, the Ewing Marion Kauffman Foundation and the Social Sciences and Humanities Research Council (SSHRC) of Canada is gratefully acknowledged. All errors remain our own.

At least one co-author has disclosed additional relationships of potential relevance for this research. Further information is available online at http://www.nber.org/papers/w31788

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

AI Adoption in America: Who, What, and Where

Kristina McElheran ⓡ J. Frank Li ⓡ Erik Brynjolfsson ⓡ Zachary Kroff ⓡ Emin Dinlersoz ⓡ Lucia S. Foster ⓡ Nikolas Zolas

NBER Working Paper No. 31788
October 2023
JEL No. M15,O3

ABSTRACT

We study the early adoption and diffusion of five AI-related technologies (automated-guided vehicles, machine learning, machine vision, natural language processing, and voice recognition) as documented in the 2018 Annual Business Survey of 850,000 firms across the United States. We find that fewer than 6% of firms used any of the AI-related technologies we measure, though most very large firms reported at least some AI use. Weighted by employment, average adoption was just over 18%. AI use in production, while varying considerably by industry, nevertheless was found in every sector of the economy and clustered with emerging technologies such as cloud computing and robotics. Among dynamic young firms, AI use was highest alongside more-educated, more-experienced, and younger owners, including owners motivated by bringing new ideas to market or helping the community. AI adoption was also more common alongside indicators of high-growth entrepreneurship, including venture capital funding, recent product and process innovation, and growth-oriented business strategies. Early adoption was far from evenly distributed: a handful of “superstar” cities and emerging hubs led startups’ adoption of AI. These patterns of early AI use foreshadow economic and social impacts far beyond this limited initial diffusion, with the possibility of a growing “AI divide” if early patterns persist.

Kristina McElheran
University of Toronto Scarborough
105 St. George Street
Toronto, ON M5S 3E6
k.mcelheran@utoronto.ca

ⓡ

J. Frank Li
University of British Columbia
jfrankliwork@gmail.com

ⓡ

Erik Brynjolfsson
Stanford Digital Economy Laboratory
353 Jane Stanford Way, Office 136
Stanford, CA 94305
and NBER
erik.brynjolfsson@gmail.com

ⓡ

Zachary Kroff
Center for Economic Studies
U.S. Census Bureau
4600 Silver Hill Road
Washington, DC 20233
zachary.kroff@census.gov

ⓡ

Emin Dinlersoz
Center for Economic Studies
U.S. Census Bureau
4600 Silver Hill Road
Washington, DC 20233
emin.m.dinlersoz@census.gov

ⓡ

Lucia S. Foster
US Bureau of the Census
Washington, DC 20233
lucia.s.foster@census.gov

ⓡ

Nikolas Zolas
Center for Economic Studies
U.S. Census Bureau
4600 Silver Hill Road
Washington, DC 20233
nikolas.j.zolas@census.gov

1 Introduction

Artificial intelligence (AI) may spark a revolution in business and the economy (Spulber 2011b). However, systematic evidence on who the using firms are, what technologies they are adopting, or where they are located, has been slow to accumulate. This makes it difficult to understand AI's economic and managerial implications or ground predictions in objective data. Yet, the consequences of uninformed AI investment and policy may be far-reaching (Agrawal et al. 2019) and early insights are essential (Goldfarb et al. 2023). To address this, we leverage new large-scale survey data to characterize early AI use across the U.S. economy, establishing a representative baseline and surfacing factors shaping its impact on workers, firms, and society.

Organizational contexts, in particular, are known to influence the use and outcomes of technology (Bresnahan et al. 2002), including recent advances such as data analytics (Aral et al. 2012; Wu et al. 2019; Brynjolfsson et al. 2021). Yet, today concern is growing over the contexts and consequences of rapidly-advancing innovations. Income inequality is on the rise, with emerging evidence linking this to digitization (Autor et al. 2020; Tambe et al. 2020; Barth et al. 2023). AI-related technologies, in particular, are attracting attention and debate regarding the "future of work" (Brynjolfsson 2022; Acemoglu and Restrepo 2020). Meanwhile, limited insight into organizational factors associated with AI usage limits guidance regarding specific areas of promise or concern.

To shed light on patterns of AI use among firms, we rely on a new nationally representative survey, the Annual Business Survey (ABS), conducted jointly between the US Census Bureau and the National Center for Science and Engineering Statistics (NCSES) within the National Science Foundation. The ABS solicits information on firm-level use of several advanced technologies—including those currently associated with progress in AI—over the prior year, beginning in 2018. The initial wave of the survey was sent to a large, representative sample of over 850,000 firms in the private non-farm economy, with a high response rate of 69%.[1] We link this to the Longitudinal Business Database (LBD), which provides administrative data on employment and revenue throughout the firm life cycle for a near-universe of US firms. This allows us to characterize early AI use in a sample of 474,000 firms which, once weighted appropriately, is representative of over 4 million firms nationwide. We further leverage rich founder and organizational data for a subsample of 75,000 startups to provide insight into the relationship between AI use and firm dynamics.

We first analyze *who* adopted AI, defined as firm use of at least one AI-related technology in production.[2] This directly captures application of these emerging technologies, yielding insights into their potential economic impact that are distinct from alternative measures such as invention (e.g., patents) or skill demand (e.g., job postings). Our representative statistics indicate that just under 6% of firms nationwide used AI as of 2017. Yet most very large firms (over 5,000 employees) reported at least some AI use, leading to employment-weighted adoption of 18%. Intensity varied from merely testing (1.1%) to using AI in more than a quarter of production (2.2%). More-intensive use was prevalent among 25-30% of the largest firms, attesting to its skewness.

Early AI adoption varied significantly by industry. Manufacturing and information were leading sectors, with adoption rates of roughly 12%, each. Yet the breadth of AI's potential is evident in its presence across every sector of the US economy, as well as in the diversity of its applications.

Our ability to observe *what* was adopted, in terms of specific business technologies (including, but not limited to, AI), reveals key interdependencies. Firms using AI were more likely to rely extensively on both digitized information and cloud computing, suggesting important complementarities with other "enabling" technologies (Bresnahan and Trajtenberg 1995; Kapoor and Teece 2021) in AI's diffusion. Also, while typically studied in isolation, we find combined use of AI and robotics: the majority of firms using robotics in production also used AI. Such clustered use of emerging technologies points to mutually-reinforcing technical and process innovation, which tend to promote broader economic impact over time (Rosenberg 1963; Petralia 2020; Goldfarb et al. 2023).

Despite our single year of survey data, we gain insights into firm dynamics, as well. Visibility to the firm life cycle reveals that, conditional on industry and size, younger firms were more likely to use AI. Further, given the central role of young firms in innovation and economic dynamism (Spulber 2011a; Acemoglu et al. 2018; Haltiwanger et al. 2013; Decker et al. 2017, 2020), we expect that patterns of AI adoption among growth- and innovation-oriented startups will be essential to understanding the rate and direction of AI use in the economy, moving forward.

We thus delve deeper into a large subsample of roughly 75,000 startups (5 years or younger) for which we have unusually detailed data on owner characteristics and motivations, startup financing, and innovation and business strategies. Across a range of specifications, we find that AI-using startups tended to have younger, yet more highly–educated and more-experienced, leaders. Founders motivated to bring new ideas to market and help the community were more prevalent among AI-users than those pursuing so-called "lifestyle" entrepreneurship (e.g., emphasizing flexible hours or work–family balance). Key markers of high-growth entrepreneurship, specifically venture capital (VC) funding, high initial capitalization, and reliance on formal intellectual property (IP) protection, also predict AI presence. Recent innovation in products or processes (particularly the latter) were further associated with startups' use of AI, as were growth-oriented business strategies.

We argue that these differences—many determined at founding (Guzman and Stern 2015, 2016) or suggesting distinct strategic commitments (Ghemawat 1991; McElheran et al. 2019; Li 2023)—are fundamental to understanding AI's trajectory. Further, while markers of high growth potential (e.g., Guzman and Stern 2020) predict AI use, conditioning on these factors reduces but does not eliminate a significant association between AI use and revenue growth. This holds for later revenue growth, too, in a broader sample including older firms. While establishing causality is beyond the scope of our study, these patterns point to important linkages between AI use and firm performance.

Finally, while the spread of AI use across the country remains in its initial phase, the potential for an "AI divide" across regions and cities is attracting concern. Investigating *where* early AI use was located among startups,[3] we find considerable concentration. The expected dominance of California's Silicon Valley and Bay Area are apparent. Yet the share of startups using AI was also high in the metro areas surrounding Nashville, San Antonio, Las Vegas, New Orleans, San Diego, and Tampa. Once we weight by employment to account for worker exposure to AI, Riverside, Louisville, Columbus, Austin, and Atlanta also show highly concentrated AI use. Widening the lens to include older single-unit firms, we observe geographic dispersion beyond key "pioneer" locations (Bloom et al. 2021). Also, distinguishing employee exposure (regardless of specific skills) from that of firms may be critical for assessing the "future of work" across these geographies. Specifically, our approach flags a number of emerging hubs for AI application that are distinct from those associated with early AI development and commercialization (e.g., Muro and Liu 2021; Bessen et al. 2021).

Our snapshot of early AI use constitutes the most extensive and detailed evidence on actual AI usage, to date. Our main objective is to document—rather than explain—AI's baseline economic diffusion and signpost its potential trajectory. We provide a number of novel stylized facts with a breadth of implications. These may be collapsed, however, into five overarching themes. First is the low yet skewed early use of AI, suggesting frictions and scale advantages in adoption and adaptation by US firms. Yet, second, we observe indications of "GPT-ness"—i.e., broad economic reach and growth potential, consistent with transformative "general-purpose technologies," (Bresnahan and Trajtenberg 1995; Trajtenberg 2018; Cockburn et al. 2019; Goldfarb et al. 2023). While initially concentrated among the largest firms, AI use spans all sectors of the economy and appears alongside other high-potential technologies such as cloud computing and robotics, as well as correlating strongly with process innovation. Third, AI use is associated with markers of high-growth entrepreneurship, as well as actual revenue growth. Our unusually-detailed analysis of over 75,000 startups reveals a range of organizational and geographic factors that have already begun to shape AI's deployment in ways that have been unobserved, far less appreciated. Among these, the importance of firm strategies related to innovation and growth stands out as a fourth key finding. Fifth is geographic concentration, including some novel "hubs" based on use in production (versus invention or skill demand) and rankings that shift when we consider employee exposure to AI.

These patterns of early use signpost where AI is likely to have the most (or at least the earliest-observable) impact. For example, the connection between early AI adoption and process innovation points to a classic friction between a novel technology's potential and its ultimate organizational and economic impacts (Bresnahan and Greenstein 1996; McElheran 2015; Feigenbaum and Gross 2021). Yet the separate importance of formal intellectual property (IP) protection suggests that early adopters are investing in protecting their innovative efforts. This, along with growth-oriented business strategies, have important implications for how firms intend to compete using AI (Iansiti and Lakhani 2020). More speculatively, higher AI use among firms with prosocial founding motivations may provide guardrails against unethical or biased applications (Cowgill and Tucker 2019), while early concentration among large incumbents may augment existing "superstar" firm dynamics (Autor et al. 2020; Tambe et al. 2020; Camuffo et al. 2022). Geographic concentration evokes similar concerns about an AI-fueled "digital divide" if early patterns persist.

Much work remains to unpack these and other implications of our findings. This study illustrates, however, numerous ways in which future assessments of AI's economic and social outcomes may hinge on choices over input and outcome measures, data availability, and which subpopulations contribute meaningfully to underlying data-generating processes.

This study makes a few main contributions. First, our approach allows us to characterize the early—and, by extension, the most likely later—AI adopters in the largest and most representative sample of firms to date. This approach is critical for getting the "denominator" right in estimating rates of AI diffusion and identifying likely areas of impact. Prior work has estimated AI's potential using O*NET task descriptions (e.g., Brynjolfsson et al. 2018; Felten et al. 2021; Eloundou et al. 2023), online job postings (Alekseeva et al. 2021; Acemoglu et al. 2022; Babina et al. 2022; Bessen et al. 2021; Goldfarb et al. 2023), and patents (Webb 2019; Hong et al. 2023; Miric et al. 2023). However, direct firm-level measures of AI use in production are rare, as are rich organizational "intangibles" (e.g., owner motivation or recent process innovation)—particularly at scale. Privately held, younger, and smaller firms—essential to contextualizing the phenomenon in the overall economy—are also typically lacking. The hurdles producing such measurement gaps are non-trivial (McElheran 2018; Miric et al. 2023), yet not insurmountable—as we show here. This new survey may serve as a guide for follow-on work across other country and research contexts.[4]

Second, tying AI use to owner characteristics and motivations in a large subsample of young firms adds to the nascent literature on entrepreneurial motivation (Cassar and Meier 2018; Gans et al. 2019; Shah et al. 2019; Guzman et al. 2020; Ganguli et al. 2021). We document novel linkages between startup leadership, high-growth entrepreneurship, and adoption of frontier technologies.

Third, our findings contribute updated insights to an important literature on the drivers of information technology (IT) adoption, including organizational complements to IT use (e.g., Ichniowski et al. 1995; Brynjolfsson and Hitt 2000; Black and Lynch 2001; Bresnahan et al. 2002; Bloom et al. 2012). More-recent work in this vein has emphasized innovation and business strategy as complements to data and analytics (Wu et al. 2020; Brynjolfsson et al. 2021), while documenting the unevenness with which the "digital age" has arrived (e.g., Brynjolfsson and McElheran 2016). Identifying the correlates of AI adoption is an essential step in this agenda. Further, our focus on elusive high-growth startups, as well as the large number of typically-unobserved organizational details we explore, both add new dimensions to this expanding and timely body of work.

Finally, this study speaks to a literature on the geography of innovation and technology diffusion, dating back to Griliches (1957), which has gained momentum with the rise of digital technologies (e.g., Forman et al. 2012; Tambe 2014; Kerr and Robert-Nicoud 2020; Bloom et al. 2021). While certain familiar patterns of concentration emerge with respect to AI, we find nuances in the relationship between urbanization and AI adoption. Moreover, we emphasize that employment-weighted geographic trends are indicative of how workers, not solely firms or those with AI-specific skills, may be exposed to AI—a distinct perspective enabled by our comprehensive administrative data.

Our analysis shows that AI use as of 2017 was low, yet concentrated in already-leading firms and geographies—raising concern about the further rise of "superstar" firms and places. That said, its penetration across all industrial sectors and association with other emerging technologies—along with its presence among the high-growth startups essential to economic dynamism—point to significant and potentially broad-based economic and organizational impact in the future. Our findings establish an early baseline of actual deployment of AI in US production. They also surface a number of novel insights into the industrial, technological, and organizational context of AI use to inform how we anticipate and navigate AI's pitfalls and potential, moving forward.

2 Motivation and Prior Work

Our framework for exploring the prevalence and implications of AI use builds on Zolas et al. (2020), which reported on the use of advanced business technologies, more broadly. Here, we focus specifically on AI and the technological and organizational context surrounding its use.

2.1 Technology Diffusion and Its Implications

A longstanding literature seeks to understand patterns of technology diffusion and their economic and social outcomes (e.g., Griliches 1957). A key motivation is the relationship between technology use and productivity (e.g., David 1990; Brynjolfsson and Hitt 1996), as well as the role of "general-purpose technologies" (GPTs) in innovation and technical progress (Bresnahan and Trajtenberg 1995). Recent digital advances are gaining attention, including "big data," analytics, and data-driven decision making (Brynjolfsson and McElheran 2016, 2019; Tambe 2014; Wu et al. 2019; Brynjolfsson et al. 2021). Yet, the speed of technological change has raised questions about how well intuitions based on prior technologies apply to AI (e.g., Agrawal et al. 2019; Trajtenberg 2018; Acemoglu and Restrepo 2020; Goldfarb et al. 2023).

Recent research into AI's potential impacts has relied on O*NET task descriptions (Brynjolfsson et al. 2018; Felten et al. 2021; Eloundou et al. 2023), online job descriptions (Alekseeva et al. 2021; Acemoglu et al. 2022; Babina et al. 2022; Goldfarb et al. 2023), or combining the latter with patents and/or research publications (Webb 2019; Bessen et al. 2021; Hong et al. 2023).

Direct firm-level measures of AI use are rare, particularly outside of Europe (where samples tend to be smaller—see, e.g., Czarnitzki et al. 2023; Hoffreumon et al. 2023). Studies of AI in other countries, such as Canada and China, have relied on yet different measures (Alexopoulos and Cohen 2018; Lu et al. 2023; Beraja et al. 2023). In addition, privately held, younger, and smaller firms—essential to contextualizing the phenomenon in many economies[5]—are systematically underrepresented, as are their organizational details. Our study makes a contribution not only by measuring actual use of AI in production, but also in terms of the scale of the data set, its representation of the diversity of US firms, and the detailed insights into the industrial, technological and organizational contexts (particularly for startups, discussed below) surrounding AI use.

Earlier IT diffusion research established the importance of commonly observed firm characteristics such as industry, size, and location (e.g., Forman et al. 2005). We characterize AI use along these familiar dimensions, though with unusual granularity when it comes to industry. We also disentangle often-confounded factors such as age and use of other advanced technologies. The former is often confounded with size, despite being central to firm dynamics (Kueng et al. 2014; Haltiwanger et al. 2013; Criscuolo et al. 2014). Our exploration into the latter, which focuses on digitization and cloud computing, speaks to a rising focus on "enabling" technologies (Rathje and Katila 2021; Gambardella et al. 2021; Kapoor and Teece 2021), including recent findings linking data, cloud computing, and AI (Goldfarb et al. 2023; Lu et al. 2023; Bessen et al. 2022). Identifying overlaps among emerging technologies has been central to GPT research seeking to understand the reinforcing patterns of technical and process innovation that may enable a "take off" in economic growth (Rosenberg 1963). Studying AI in isolation from these other high-potential innovations could obscure important interdependencies shaping its diffusion and impact.

Similarly, while AI and robotics are frequently discussed together as automation technologies primed to transform the global workforce (Raj and Seamans 2019), they are typically studied in isolation. Robots primarily automate physical tasks (Dixon et al. 2021), whereas AI automates cognitive ones (Agrawal et al. 2019). Yet robotics is poised to increasingly rely on AI for greater autonomy and reprogrammability. Understanding their co-occurrence may be key to understanding how AI is mobilized in physical production, with attendant firm and labor-market outcomes.

2.2 Entrepreneurship and Technology Use

A key advantage of our research setting is detailed organizational data for a large sample of non-public, primarily young and small firms. Younger firms, in particular, are often missing from standard data sets, despite being vital to economic dynamism and innovation (Spulber 2011a; Chava et al. 2013; Decker et al. 2017, 2020; Haltiwanger et al. 2013, 2016; Acemoglu et al. 2018; Guzman and Stern 2020). We examine AI use among startups in detail, focusing first on characteristics that are largely fixed at founding and, as we establish empirically, associated with revenue growth early in the firm life cycle. Our overarching motivation for exploring the relationships to follow is that, if early AI use is concentrated among high-potential firms, we might expect greater economic and social impacts of AI as these firms scale. Conversely, if AI-related economic gains remain elusive (Brynjolfsson et al. 2021), our findings may point to key frictions or measurement challenges for future study. Finally, if we can identify a relationship between AI use and later firm growth, this may indicate where causal links between AI use and firm performance will ultimately be found.

2.2.1 Owner Characteristics: Education, Experience, and Age

We leverage a growing body of evidence that certain early organizational markers can indicate a firm's growth potential and reliance on innovation in its strategy (e.g., Guzman and Stern 2020; Botelho et al. 2021). Such high-potential entrepreneurship has further been linked to a number of founder characteristics and motivations (Levine and Rubinstein 2018; Gans et al. 2019; Guzman et al. 2020). Intuitively, founders with advanced degrees are likely to lead firms that are more technology-focused or managed using more-advanced technology. An entrepreneur's prior experience is likely to matter, as well. Serial entrepreneurship, in particular (e.g., Lafontaine and Shaw 2016), may be correlated with a better assessment and exploitation of the business opportunities advanced technologies offer, or reflect adaptability to new changing technologies. We anticipate that advanced formal education and experience will be associated with greater AI use, all else equal.

With respect to age, Azoulay et al. (2020) find the mean founder age for the fastest growing new ventures to be 45, which is older than commonly understood — particularly in technology-focused sectors. Yet it is consistent with tension between "vintage-specific human capital" (younger workers have better facility with more-recent technology – see Chari and Hopenhayn 1991) and experience or other intangible skills that take time to accumulate. When the latter complements digitization, this creates the best "fit" (and associated compensation) for mid-career workers at IT-intensive firms compared to *both* older and younger employees (Barth et al. 2023). However, prior work has not explored how this tension nets out in the context of AI adoption, whose newness may require even more familiarity with the most-recent technology, particularly at the top of the organization.

2.2.2 Founder Motivation

Next, we expect that aspirations of entrepreneurs will shape the rate and direction of AI's impact, at least among startups. The smaller older firms that lag in advanced technology use (Zolas et al. 2020) may have been founded to pursue objectives other than growth, lacking the use cases, complementary inputs, and scale to justify AI adoption. The literature has delineated "lifestyle" from "high-growth" entrepreneurship (e.g., Hurst and Pugsley 2011). We anticipate that lifestyle-focused owners will be less likely to make the technology and organizational investments required to deploy AI in a meaningful way.

Recent work also distinguishes between founders motivated by prosocial aspirations and/or workplace values and those motivated primarily by earning potential (Cassar and Meier 2018; Guzman et al. 2020; Shah et al. 2019; Ganguli et al. 2021). While we cannot observe specific applications of AI, we expect founder motivation may foreshadow the potential for ethical or prosocial applications, versus ones with more-adverse consequences (e.g., Cowgill and Tucker 2019).

2.2.3 Startup Conditions

Firm trajectories have further been linked to aspects of startup financing. Selection by—and synergies between—VCs and entrepreneurs have been linked to firm growth.[6] A relationship between venture capital, high-growth entrepreneurship, and AI use may therefore arise for a variety of reasons. One channel could be selection: VCs may be able to better identify startups capable of developing or leveraging new AI technologies. Another could be treatment: the type of products and processes venture capitalists encourage may rely more significantly on AI use.

Higher initial capitalization may also reflect private information about firm growth potential. We thus also explore its separate relationship to AI use. Care is required, however, as capital intensity, even within narrowly defined industries, may reflect production strategy (McElheran et al. 2019) as much as firm "quality," and is positively associated with the presence of other advanced technologies in manufacturing plants (Dinlersoz and Wolf 2018). We leverage direct insights into dimensions of innovation and business strategy (rarely observed at scale) to disentangle the source and level of financing from growth or "quality" as factors in AI adoption.

2.3 Startup Innovation and Business Strategies as AI Complements

Many of the organizational factors discussed above are relatively fixed early in the firm life cycle. However, we are also interested in the relationship between AI use and startups' innovation and business strategies, recognizing that they may be co-determined. A large and influential literature explores complementarities between organizational features and technology use (e.g., Ichniowski et al. 1995; Brynjolfsson and Hitt 2000; Bresnahan et al. 2002; Bartel et al. 2007; Bloom et al. 2012; Aral et al. 2012). This literature, which focuses predominantly on large incumbents, typically argues that organizational characteristics should be considered quasi-fixed in comparison to IT investment (e.g., Tambe et al. 2012), and takes as a key test of complementarity the organization-technology correlations of the type we explore, here.[7]

Within this large and growing stream of research, the notion that commitments (Ghemawat 1991) to certain innovation and business strategies might promote or complement technology use (rather than vice versa) has received relatively scant attention. Concerning innovation, process versus product innovation is an important distinction—indeed, a tension—in firm dynamics and innovation strategy (Cohen and Klepper 1996). Whether and how this distinction matters with respect to frontier IT use, however, is less clear. On the one hand, "informal" process innovation may be critical for successful adoption of any new technology (Bresnahan and Greenstein 1996; McElheran 2015; Feigenbaum and Gross 2021). On the other, formal intellectual property (IP) protection may signal commitments to innovation and future growth (Guzman and Stern 2020; Dinlersoz et al. 2021, 2023) that promote both AI use and economic impact. Understanding nuances in this relationship may be essential to unpacking adoption and growth dynamics related to AI.

Finally, recent work indicates that dimensions of a firm's business strategy will shape both the application and performance of new technologies (Wu et al. 2020; Brynjolfsson et al. 2021; Li 2023; McElheran and Jin 2023). We extend this line of inquiry to understand AI use among startups with different business strategies, focusing on the distinction between growth-oriented strategies versus others, such as low-cost or "niche" positions (e.g., Porter 1980).

2.4 Geography

Research into both entrepreneurship and technology diffusion have long emphasized the role of geography in patterns of technology use, location choice, and economic outcomes (e.g. Forman et al. 2012, 2016; Delgado et al. 2010; Alcácer and Chung 2014). Recent work has pointed to growing concentration, both of new digital technologies (Tambe 2014; Bloom et al. 2021), and of startup activity (Guzman and Stern 2015), particularly in California (Guzman 2019). VC- and innovation-related spillovers (Jaffe et al. 1993; Samila and Sorenson 2010) may further contribute to the emergence of "hubs" for startups leveraging specific technologies, though this is often difficult to disentangle from other agglomeration economies (Kerr and Robert-Nicoud 2020). While AI is still in the early stages of its geographic diffusion (Muro and Liu 2021; Hong et al. 2023), "pioneering" tech locations have been found to retain a considerable share of the relevant employment (Bloom et al. 2021), and distance from "hotspots" of AI invention has slowed its adoption for certain types of firms (Bessen et al. 2021). Thus, AI's future path will likely depend on its early geographic dispersion. We do not seek to pin down specific drivers, focusing instead on describing the dispersion of AI use in production and its implications for workers (regardless of their specific skills) That said, uncovering geographic patterns in AI adoption is a key objective of our analysis.

3 Data and New Descriptive "Facts"

This study provides the first in-depth look at AI adoption from the Annual Business Survey (ABS). Introduced in 2018, the ABS consolidates three earlier data collections capturing important features of US firms. New designed-for-purpose questions on technology use were added to existing questions that provide, for relevant firms, ownership and owner characteristics, startup financing, intellectual property use, and other details of innovation and business strategy.[8] We describe data construction and novel descriptive statistics here, beginning with our baseline sample. For additional details, see Data Appendix B.

3.1 Main Sample Descriptions

The 2018 ABS was sent to 850,000 firms nationwide. Approximately 583,000 firms responded to the survey, 573,000 of which were linked to the Longitudinal Business Database (LBD). The

LBD is curated by Census to provide a comprehensive panel of microdata on employment, revenue, and payroll for the private non-farm economy (Chow et al. 2021). The LBD tracks firms from their birth as an employer to death, accounting for mergers and acquisitions. This provides the basis for our systematic exploration of early- and later-stage growth patterns for firms in the ABS.

Leveraging the LBD mitigates concerns that we could overlook dynamic entrepreneurial firms, which are often missing from standard data sets comprised of larger incumbents (e.g., Compustat). Conversely, it ensures we do not omit older firms that might be systematically excluded if sampling depended on the presence of advanced technology (such as participating on a digital platform). We call this sample of 573,000 firms the ABS–LBD linked sample (column 1, Table 1).

Establishing the correct "denominator" is a key objective of our study and advantage of our access to administrative data. This requires additional sample restrictions in certain instances. For many of our core stylized facts related to AI use, we restrict attention to the 447,000 firms for which the AI-related questions are minimally complete (basically not completely missing or unknown to the respondent—see Data Appendix B for details). Descriptions of average firm size and age for this "baseline" sample are presented in column 2 of Table 1. We weight the firms in our samples to better represent the firm size, age, and industry distributions of the population of employer businesses in the country-wide Business Dynamics Statistics (BDS).[9] The vast majority of US firms are very small,[10] which is reflected in our baseline sample. The average firm, here, has roughly 62 employees, but only 21 when weighted. Average firm age is around 16 years (regardless of weighting). Weighting makes our baseline sample representative of over 4 million firms across the US economy (column 2, Table 1).

Investigating the organizational context, linkages to revenue growth, and geography of AI use impose other sample restrictions. These vary by analysis and are discussed as they arise, below.

3.2 New Measures for the Digital Age: Digitization, Cloud Computing, and Advanced Business Technologies

The technology module in the 2018 ABS contains three new and connected questions on digitization and technology use. The first queries firms' reliance on data, widely regarded as a key input to more advanced uses of digital technologies (Brynjolfsson and McElheran 2016, 2019). Over 65% of the ABS-LBD linked sample reported having at least one type of information in digital format.

A second oft-cited enabler is sufficient computing power to store and analyze massive quantities of data. Thus, the second question asks about the extent to which firms rely on cloud services, which have shifted the cost structure and speed of access to IT resources for many firms (Armbrust et al. 2010; Jin and McElheran 2017; Goldfarb et al. 2023; Lu et al. 2023). In the same representative sample, 43% of US firms purchased cloud services for at least one IT function.

The third question, and primary focus of this paper, asks about the use of various advanced business technologies in producing goods or services, including intensity of use. Specific definitions are provided in Table 2 and include five commonly associated with AI (bolded): machine learning, machine vision, natural language processing, voice recognition software, and automated guided vehicles.[11] Tabulated responses to these five items, conditional on some use of the technology in question, are reported in Figure 1. Here, it is useful to keep in mind that only 2% of firms in the overall unweighted ABS-LBD linked sample reported any use of AI at all.[12] Thus, the y-axis represents the number (not share) of firms. Note the considerable dispersion across both specific technologies and intensities of use. To our knowledge, this is the first study to capture this detailed variation in firm-level AI use.

Also reported on Figure 1, adoption rates of specific AI-related technologies were below 3%, across the board—ranging from a high of 2.9% in the case of machine learning, to a low of 0.8% for automated guided vehicles. Testing of each technology—presumably a leading indicator of use in production—was even lower (less than 1% each), suggesting that these levels were not poised to change dramatically in the wake of our study.

3.3 Early AI Use: Low and Skewed

For our main analyses, we collapse different technologies and intensities of use into a single binary indicator of *any* use of AI in production. According to this definition, the share of firms in the baseline sample that adopted at least one AI-related technology as of 2017 was 5.8%. The share of firms testing at least one of these technologies (but not yet using it in production) was 1.1% (see Figure 1).

This level of AI diffusion is significantly lower than that reported elsewhere, such as the European Commission survey of AI (Kazakova et al. 2020) or other private surveys by McKinsey (Chui and Malhotra 2018), Deloitte (Deloitte 2018), or PwC (PwC 2019). However, these surveys do not claim to be representative, oversampling larger, often publicly traded companies. In contrast, our baseline sample includes the many small firms for which AI adoption is quite limited—yet whose inclusion allows us to estimate the correct "denominator" for country-wide AI adoption rates.

While these low average rates are informative—and useful for counteracting potentially excessive "exuberance" related to AI—they paint an incomplete and potentially understated picture of AI use in America. As for many advanced business technologies (e.g., Zolas et al. 2020), early AI use appears highly skewed and concentrated among larger firms. Figure 2a reports AI use by firm size for different intensities.[13] Note that the majority of firms with more than 5,000 employees reported at least some use of AI. Over a quarter of the largest firms reported using AI intensively.

This skewness in both AI presence and intensity underscores the potential impact of AI not only on firms, but also on employees at these firms. The concentration of AI in the right tail of the firm size distribution means that worker exposure (regardless of their specific skills) to AI may be considerably higher. Returning to Figure 1, employment weighting (in parentheses) suggests machine learning prevalence of 8.5%, and employment-weighted usage of *any* AI technology of roughly 18.2%. We leverage employment weighting to establish other stylized facts, below, focusing on firm-weighted statistics, next.

3.4 *Who* Adopts AI: Industry, Size, and Age

Our rich administrative data grants insights into key characteristics of early AI users—the *who* of interest in this study. Figure 2b reports AI use and intensity by sector. Manufacturing- and information-sector firms led in early AI adoption, with roughly 12% extensive-margin adoption rates, each. Both also led in intensity of use, with health-related firms close behind, followed next by professional services. Lagging sectors included construction and retail trade, at roughly 4%, each. Finance, insurance, and real estate ("FIRE") sectors had adoption rates below 6% with limited intensity, despite ranking highly in digitization—see, e.g., Figure 2 in Zolas et al. (2020).

Table 3 zooms in on several non-manufacturing industries (at the 4-digit NAICS level) where AI was most prevalent. Medical and diagnostic laboratories had a striking uptake above 23%. Other high-AI industries tended to be in "high tech," such as software publishing (nearly 16%) and computer systems design (14%), though with exceptions, such as physicians' offices (19%). Data processing and related services, which includes cloud services, also reported AI use around 14%.

This is notable, as many of these represent industries where we expect development of *applications* of AI for use in other sectors of the economy (e.g., Bresnahan 2023).

We control for this significant industry-based heterogeneity at the 6-digit NAICS level (a degree of granularity difficult to achieve in standard data sets)[14] to highlight other patterns of interest. Table 4 reports the likelihood of using AI by size and age, including state-by-industry controls. Conditional correlations show AI usage increasing monotonically with size. In column 1, firms in the top percentile of their industry's employment distribution were nearly 9 percentage points (pp) more likely to use AI compared to below-median firms. In contrast, we find AI presence decreasing almost monotonically with age, though magnitudes are smaller. The oldest firms were 1–2pp less likely to use AI than firms younger than their industry's median age.[15]

3.5 *What* Gets Adopted: Technological Context and Interdependencies

Column 2 of Table 4 sheds new light on *what* technologies get adopted, adding indicators for digitization and cloud computing. Controlling for the above observables, digitization is predictive of AI use: firms that reported having at least one type of information in digital format were 2pp more likely to report some AI use. Firms that reported purchasing cloud services for at least one IT function were 5pp more likely to use AI. Further, these margins of technology use appear interconnected, per Figure 3, which shows a Sankey Diagram of firms that used one or more of each type of technology. Firms that reported early AI use were most likely to perform their data analysis using cloud-based IT services. These same firms were also most likely to store their production information digitally.

We further observe variation by industry in terms of *what specific* AI-related technology (among those in Table 2) was most prevalent. Reported in Appendix Table A1, while more-generic "machine learning" was highly prevalent in computer systems and data processing industries, industries not always considered "high-tech" led in the use of specific technologies. For example, automated guided vehicles (AGV), while used by fewer than 1 percent of firms nationwide, were present at roughly 9 times that rate in support activities for crop production in the agriculture sector. Similarly, machine vision was most prevalent in technical and trade schools in the education sector (nearly 5%), while voice recognition was most prevalent in healthcare settings and legal services.

This industry-based heterogeneity also hints at potential complementarities between AI and other labor-saving, capital-intensive technologies such as robotics. Robotics adoption based on ABS responses (1.3% of the ABS-LBD linked sample) was significantly lower than that for AI. However, the overlap with AI appears fairly significant, with the majority of robotics adopters (57.5%) reporting using at least one AI technology. Figure 4 provides a sectoral breakdown of the overlap of firms that use AI, robotics, or both, with the greatest overlap in the manufacturing, wholesale, and education sectors.[16]

3.6 AI in Entrepreneurship: Owner Characteristics, Startup Conditions, and Firm Innovation and Business Strategies

In addition to novel technology measures, the ABS contains, for a large subsample, detailed information on the firm's owners, their reasons for owning the business, financing at startup or acquisition, intellectual property strategy, and related innovation and business strategies—key issues in entrepreneurship research.[17] Given this advantage of the ABS and the importance of young firms in job creation and innovation (e.g., Spulber 2011a; Haltiwanger et al. 2013; Botelho et al. 2021), we narrow our focus to analyze, in depth, the roughly 75,000 startups (5 years old or younger) for which we have owner and revenue data. Firms in this "startup sample" (column 4, Table 1) are just under 3 years old, on average, with around 10 employees and representative on a firm-weighted basis of 740,000 firms. Here, we describe our measures and report unconditional means, both across this sample as well as by whether or not the firm used AI, in Table 5.

3.6.1 Owner Education, Experience, Age, and Aspirations

The ABS asks roughly 20 questions about up to four business owners; we focus on the primary owner (based on ownership share). See survey details in Data Appendix B.

Panel A of Table 5 focuses on characteristics of the primary owner. Starting with education, while 22.4% of all primary owners in our startup sample reported holding an advanced degree (master's, doctoral or professional—see column 1), this rises to 31.5% among AI-using startups (column 2). Owners of AI-using startups were relatively more likely to have owned a prior business (43.8%) compared to those without AI (41.2%, see column 3). Owners of AI-using startups also tended to be slightly younger, as well.[18] We delve further into these unconditional relationships in a regression framework, below.

Panel B in Table 5 describes the possible "very important" reasons that owners (primarily founders) may report for owning the business. The biggest differences between AI-using startups and the rest (see column 4) were "Help and/or become more involved in my community" (difference of 6.6pp) and "Best avenue for my ideas/goods/services" (difference of 5.1pp). The share of owners in AI-using firms reporting these reasons were 30.3% and 61.8%, respectively, compared to 23.7% and 56.7% of startups not using AI. The prevalence of these growth-oriented and prosocial motivations contrasts with the lower likelihood of "lifestyle reasons" for entrepreneurship among AI users. "Flexible hours" or "Balance work and family" were, respectively, "very important" for 52.8% and 55.1% of owners of AI-using firms. This contrasts with 54.5% and 56.4% of all primary owners in the sample. Other reasons are tabulated for completeness in Table 5; however, a binary indicator of lifestyle-focused, versus other motivations turns out to be the most intuitive and empirically important in a richly-specified regression framework (see below), though the unconditional means are not strikingly different at 0.8pp.

3.6.2 Startup Financing

The ABS further contains rich information on the source and size of capitalization at startup or acquisition, described in panel C of Table 5. For VC funding, which was uncommon across the sample, we observe a sizable difference in means: 2.9% of AI-using startups were funded by venture capital (column 2), compared to only 0.9% of all startups in the sample (column 1). AI-using startups also reported larger capitalization, on average. Combining rows in this panel, 46.5% reported startup capital of more than $25,000 compared to only 41.4% of all startups; 3.8% reached or exceeded $1 million, vs. 2.4%, on average.

3.6.3 Innovation and Business Strategies

The greatest unconditional differences between AI-using and other startups centers on innovation and business strategies, as shown column 4 of panel D, Table 5. Process and product innovation within the last 3 years (2015–2017) are elicited, separately on the ABS. Process innovation, which includes innovation in methods of manufacturing, logistics or distribution, or in supporting activities, while reported by only 19.9% of startups, overall, was substantially more prevalent among AI users at 39.3%. Product innovation (goods and/or services) was both more common overall (53.6%) and more prevalent among AI users (66.3%).

AI-using startups relied on formal IP at higher rates, as well. The prevalence of patents owned or pending among such firms was 5.2%, versus 2.1%, overall. IP was reported to be "very important" at a rate of 40.4% among AI-users, versus only 18.1% among those not using AI in production.

We collapse the fourteen business strategies measured in the ABS (see Appendix B) into a single indicator of how "growth-oriented" the firm's business strategy is, based on whether it reported that introducing new goods or services, reaching new customer groups, or opening up new domestic or export markets were "very important" AI-using startups espoused a growth-oriented strategy at a rate of 77.5%, compared to 68.9% of non-AI-users.

3.7 *Where*: The Geography of AI-Using Startups

The last set of descriptives in Table 5 speaks to the *where* of our study. In panel D, we take as our unit of aggregation Core-Based Statistical Areas (CBSAs).[19] While a sizeable share of startups in our sample (22.8%) are in very small CBSAs (fewer than 50,000 people), over 50% are in metro areas of 1 million or more, with 55.1% of AI-using startups located in a large CBSA.

To better understand the prevalence of AI use in production across these areas, we next expand our focus to include all young single-unit firms (regardless of owner/organizational data). In addition to being straightforward to situate geographically (compared to multi-unit entities), single-unit firms accounted for approximately 97% of all US firms and 75% of establishments, respectively, per the Census Bureau's Business Dynamics Statistics for 2017. Further, while our main analysis of geography retains the focus on startups from above, we also exploit the unusually good data coverage to map all AI-using single-unit firms, regardless of age, that we can observe (see below).

We assess AI prevalence within each CBSA, first, by the share of single-unit startups that report using AI in production.[20] The largest CBSAs are ranked by share and mapped in Figure 5a.[21] Bubble sizes represent the number of AI-using startups present in the CBSA, while the color gradient indicates the percentile rank of the CBSA in terms of AI presence—lighter colors correspond to greater prevalence of AI use among young single-unit firms in that CBSA.

Some high-level patterns are apparent. Areas known for pioneering technologies, such as Silicon Valley in California and Research Triangle in North Carolina,[22] show high concentrations of AI use among startups as shown by the lighter colors. Yet hubs for AI use are apparent throughout the South and the West, as well. The share of startups using AI was high in the metro areas surrounding Nashville, San Antonio, Las Vegas, New Orleans, San Diego, and Tampa.

The Northeast and Midwest display lower AI *intensity* as share of the firm population (represented by darker colors). However, larger metro areas have a larger number of firms in general. Thus, large metro areas such as New York City, Los Angeles, Chicago, and Washington D.C. display prominently in terms of bubble size (i.e., have many single-unit startups that use AI), while *within* these areas, AI usage is not particularly intense.

This view of AI geography among startups shifts somewhat when we use employment-weighted adoption rates across CBSAs, as shown in Figure 5b. This assessment, based on the share of employment among single-unit startups accounted for by AI-using firms, provides a view of worker exposure to AI. Note that this is regardless of firm demand for particular skills, and assumes that the use of AI in production at the firm constitutes exposure of all of its employees to AI.

Among US startups, Riverside, California, actually supercedes the Bay Area on an employment-weighted basis (see Table A2). Additional locations attain "hub status" using this approach, primarily in the Midwest and Mid-Atlantic, including Louisville, Columbus, Austin, and Atlanta. Locations in the South tend to fall, in a relative sense, as a result of employment weighting. These shifts can be attributed to a variety of factors. For instance, certain locations have a few young AI users that are relatively large compared to others in the same location, whereas in other locations AI-using single-unit firms may be numerous yet too small to make up a large share of employment. Differences in industrial composition also matter. For instance, manufacturing is a leading AI-using industry (Figure 2b), leading large AI-using manufactures clustered in the Midwest to receive greater weight, despite comprising a small fraction of their CBSA's total single-unit startups.

Table A2 provides a list of CBSAs, ordered alphabetically, alongside their share of single-unit startups with AI use and overall ranking. Firm-weighted rankings (corresponding to Figure 5a) are in column 1, while column 2 relies on employment weighting (Figure 5b). The bottom panel summarizes overall shares for micropolitan, small-, and medium-sized CBSAs. While AI usage rates are generally lower in the smaller CBSA groups, they are not too far below the corresponding rates in the largest CBSAs. We return to the role of CBSA size in AI adoption in Section 4.3.

Widening the lens to include the large number of older, single-unit AI users in our data, we perform a similar analysis for all single-unit firms across 350+ CBSA categories (all sizes) for which at least 20 single-unit AI firms exist. See Figures 6a and 6b. The results are qualitatively similar but with some interesting patterns of dispersion among the small and mid-sized CBSAs.[23]

4 Regression Analysis

4.1 Methodology

Moving beyond novel descriptive statistics, our empirical approach assesses whether the organizational features described above may help predict patterns of AI adoption, as well as how the unconditional patterns might shift in a more richly-specified regression framework. This will shed light on the relative importance of various drivers of AI use, as well as clarify the extent to which AI adoption is associated with early indicators of a young firm's economic potential.

To do this, we proceed in steps. First, we establish the extent to which these traits do, indeed, identify startups that turned out to be "high quality" in terms of being high-growth early in life. For these analyses, we measure early growth trajectories by taking the average of the annual log-difference revenue (r_t) growth rate over the first three years of a firm:

$$g = \frac{1}{3}\sum_{t=1}^{3} \ln(\frac{r_t}{r_{t-1}}). \tag{1}$$

After using this measure to assess which firm characteristics are associated with early firm growth, we fit a linear probability model to explore the extent to which the same firm features may also predict AI use.

Beyond characterizing patterns of early AI use, we are interested in the extent to which its application may be related to firm performance. To explore this connection, we examine whether AI presence had a significant association with firm growth after controlling for our rich set of firm observables. We extend this analysis by further measuring average growth over the three years 2016–2018 (later-stage growth) around the survey reference year of 2017, again using (1). Note that to understand later-stage growth we need to expand the sample to include all firms for which we have the relevant owner and revenue data (column 3, Table 1), as startups are by definition in the early stages of their life cycle (i.e., the first 5 years). This approach leverages the LBD more fully, while nevertheless introducing other sampling pressures, including potential survivor bias [24] and the exclusion of publicly-traded and much-older firms (over 20 years old—see Appendix B).

Finally, our analysis does not aim to establish a causal link between firm characteristics and adoption, nor between adoption and firm growth. These findings will be most-appropriately interpreted as suggestive of linkages between technology use and performance, with careful causal identification left as a challenge for future work.

4.2 Markers of High-Growth Entrepreneurship

Table 6a reports the conditional correlations between early revenue growth in startups and the factors discussed in Section 2. We regress the early average revenue growth rate (equation 1) on our three broad groups of firm characteristics—first on each individual group, and then on all groups together—controlling for state-by-industry (NAICS 6) and firm age. While never statistically significant, we also include controls for owner gender in all specifications.

Column 1 reports the estimated coefficients for owner characteristics. An owner holding an advanced degree or previously owning a business is associated with a higher early growth rate. Combining rows 1 and 2, a business with an owner who both holds an advanced degree and previously owned a business is associated with a 4.4 percentage-point (pp) higher early growth rate. With respect to age, we observe a monotonic relationship conditional on education and experience: startups with an owner aged 35 or older tend to have a lower early growth rate than those with an owner younger than 35 (and firms with owners 55 or older tend to exhibit even less growth than those with owners aged 35–54). Having a "lifestyle" reason as being very important for starting the business returns a negative point estimate (though not statistically significant in this model).

Looking separately at startup conditions, column 2 indicates that VC investment and larger initial capitalization are both associated with higher early revenue growth. While rare, a young "gazelle" funded by venture capital and with startup capital exceeding $1 million would have enjoyed a 30pp higher early growth rate compared to the omitted groups.

Column 3 reports the coefficients for innovation and business strategies considered together. Each of process and product innovation, intellectual property use, and a growth-oriented business strategy is conditionally associated with higher early growth. This is particularly so for patent ownership (7.7pp higher early revenue growth) and having a growth-oriented strategy (5.4pp higher). While both statistically significant, recent innovation in processes has a much stronger connection to early revenue growth than product innovation (nearly 4pp, versus just over 1pp).

Column 4 includes all of the owner and firm characteristics together, showing the robustness of the patterns in columns 1–3. In addition, the indicator of "lifestyle" reason as being "very important" for starting the business increases in magnitude and becomes significant at the 10% level. Here, controlling for owner characteristics, startup conditions, and innovation/business strategies, an owner who reported a "lifestyle" reason as being "very important" for owning the business is associated with a 0.8pp lower early growth rate. In Column 5 we explore this relationship further by replacing the "lifestyle" reason indicator with indicators for every possible reason covered in the ABS, which are not mutually exclusive and reported in detail in Table 6b. The coefficients on both "Flexible hours" and "Balance work and family" (i.e., the reasons that make up our "lifestyle" indicator) are negative and significant (column 1 of Table 6b),[25] while "Best avenue for ideas/goods/services" and "Help and/or become more involved in community" have positive and both statistically and economically significant associations with early growth.

Column 6 of Table 6a examines the same specification, restricting the sample to the four sectors with the highest usage rates of at least one AI-related technology. Restricting attention to manufacturing, information, services, and healthcare (column 5, Table 1) reveals similar patterns, except that holding an advanced degree or previously owning a business do not have statistically significant links to early growth in these already-advanced sectors. As to founding motivations (Column 2 of Table 6b), we find similar patterns, except that "Balance work and family" and "Best avenue for ideas/goods/services" are not statistically significant, while pursuing flexible hours seems even more inconsistent with early growth within these more AI-intensive industries,

4.3 Conditional Correlates of AI Adoption

We next look at how well early growth, itself, predicts AI use in 2017, as well as the extent to which this relationship may be disentangled from the factors explored in Table 6a. Controlling for industry by state, as well as age and gender of the owner, column 1 of Table 7 indicates that higher growth in the first few years of a startup's life is positively and significantly associated with AI use.[26] A 10pp increase in the average growth rate in the first three years is associated with a 0.134pp increase in the likelihood of using AI—equivalent to a 2.2% increase when measured at the mean usage rate of 6% for the startup sample.

In column 2 of Table 7, we condition on the firm characteristics linked to early growth in Table 6a, beginning with those of the primary owner. The association between AI use and early revenue growth remains largely unchanged. However, having an owner with an advanced degree is associated with a 1.04pp increase in the likelihood of using AI (a 17% increase relative to the mean). The probability of using AI is also higher if the owner previously owned another business (1.20pp). Conditional on education and experience, having an owner aged 55 and up has a significant negative coefficient (-0.92pp). Lifestyle motivations are not statistically significant in this specification.[27]

Columns 3 and 4 of Table 7 repeat this exercise, but instead focus on startup financing and location, respectively. The presence of VC funding has a particularly strong relationship to AI use, with a striking 10.7pp increase in the probability of AI use, absent other controls. Higher startup capitalization, even at the more-modest threshold of $25,000, is also associated with a higher probability of AI use, though this also appears related to the absence of other covariates (see below). The relationship between business location and AI use in column 4 suggests that startups located in large CBSAs have a 0.74pp higher probability of using AI compared to those located in a micropolitan or rural area with fewer than 50,000 people (the omitted group).

In column 5, we include all previous covariates as well as innovation and business strategy measures.[28] Process innovation, patents, and placing importance on IP protection have the strongest conditional correlation with AI use, with probability increases of 5.4pp, 3.5pp, and 6.1pp, respectively. Firms that reported product innovation or a growth-oriented business strategy were also statistically more likely to be AI users, but at smaller magnitudes and conditional on the other covariates, including early growth and the other measures (contrasting with the large unconditional differences in Table 5).

In this richly-specified model, owner education and age fade in significance, while the presence of VC funding, reliance on IP, process innovation, and patent ownership stand out as having large and statistically significant conditional linkages to AI adoption. Recall from Table 5 that process innovation and placing importance on IP are distinctly more-prevalent in this young-firm sample. Together, these two characteristics are associated in column 5 with an increase in the probability of AI use of 11.5pp, nearly twice the average adoption rate. Layering in the less-common VC funding and patent ownership increases the probability of AI use by almost 24pp, equivalent to 4 times the startup sample's average rate of AI use. Figure 7 graphs the magnitudes for each of the factors in terms of their cumulative effects relative to the baseline likelihood of AI adoption.

Also in column 5 of Table 7, the conditional correlation of AI use with urbanization becomes weaker compared to the results in column 4. Conditioning directly on firm-level factors such as beneficial startup conditions and successful innovation appears to explain the otherwise strong-looking link between AI use and the amenities associated with locating in a large CBSA.

Across columns in Table 7, early growth *per se* retains a statistically significant relationship with AI adoption. However, by column 5 the magnitude falls to less than half of that in column 1. A 10pp increase in early revenue growth, here, is associated with only a 0.058pp increase in the probability of AI use (less than 1% of the mean usage rate).

At this point, a few observations are in order. The first is that a significant amount of the correlation between early growth and AI use is attributable to firm characteristics *associated* with growth, not necessarily to growth, itself. A key advantage of our study, in fact, lies in the ability to directly account for a number of factors associated with high-growth entrepreneurship, while partialling out the effects of growth that remain unexplained by those factors. This should ease future assessments of AI's impact by signposting where its early presence was concentrated.

Second, these robust conditional correlations— including the remaining association with growth early in the life cycle— reveal the use of AI among the very firms that have historically driven dynamism in the aggregate economy. Thus, whether AI use promotes growth or whether faster-growing startups have the need and/or resources to adopt AI, these associations point to a growing role for AI in US production, moving forward.

Also of note, this very richly-specified model explains only 23% of the variation in AI use, an increase of about 10% over the initial specification in column 1. While the firm characteristics included in column 5 contribute significantly to the explanatory power of the model, a substantial portion of the variation in AI adoption still remains unexplained, underscoring the large amount of still-unobserved heterogeneity underlying AI diffusion in the US.

That said, the data on which these insights rest constitute the most extensive and detailed evidence on actual AI usage, to date. This creates opportunities within the scope of this study to both deepen and broaden our understanding of early AI adoption in the US.

Narrowing the focus, we have sufficient power and coverage to restrict our attention in column 6 of Table 7 to the four leading sectors for AI use (see Figure 2b): manufacturing, information, health care, and professional/technical services. Overall patterns remain largely consistent, with a few notable exceptions. Early growth and having an advanced degree regain predictive power, compared to column 5. Owner motivation now stands out as significant: within these key sectors, AI use was 1.58pp less likely if the firm was founded or purchased for "lifestyle" reasons. This magnitude is on par with the positive correlations between AI use and growth, advanced degree, and product innovation. The latter more than doubles in magnitude within these key sectors, yet it still trails the importance of process innovation, which is associated with a 6.14pp increased likelihood of AI use. Overall, while many patterns remain robust to this shift in focus, the exercise highlights the importance of industry context in AI adoption.

Conversely, we widen the lens in column 7 to include the somewhat-older firms for which we have both owner and revenue data, recognizing the sample limitations imposed by data availability (see column 3, Table 1 and Data Appendix B). Extending the sample back to 1997, we can include not only the average revenue growth for the initial years of the firm, but also the average revenue growth rate from the 3 most recent years for which revenue information is available, 2016–2018. (While we report the complete set of results for this sample in the online appendix Table A4, we report the fullest specification in Column 7 of Table 7.) In this sample of around 209,000 firms, the coefficient on early revenue growth weakens and loses significance with the inclusion of owner characteristics, startup conditions, and innovation and business strategies. Instead, the coefficient on later-stage revenue growth is statistically significant and with a magnitude that, while small (a 10pp increase in later stage revenue would be associated with a 0.041pp increase in the likelihood of AI use) is on par with factors such as prior founding experience and greater than the effects associated with product innovation or growth-oriented business strategy.

This later period of revenue growth more likely falls into the post-adoption period for many (largely older) firms, suggesting that the use of AI itself could have contributed to the performance of the firms that put it to use. A few caveats are in order, however. First, this will largely hold for very early adopters. Based on the patterns we observe in 2017, these are likely to be larger firms, even among the privately-held firms examined in depth, here. In addition, we do not observe adoption year, only whether or not AI was being used as of 2017; thus, we cannot establish a credible causal link. That said, it is interesting to observe the statistically significant relationship between later-stage revenue growth and AI use in such a large sample, while seeing the coefficient on early revenue growth lose significance once we account for its links to owner characteristics, startup conditions, and innovation and business strategies.

5 Summary of Key Findings and Discussion

This paper captures a rich snapshot of the *who*, *what*, and *where* of early AI use across the US economy. It leverages the largest and most-detailed data collection, to date, concerning firm use of AI in production, based on the inaugural wave of the Annual Business Survey and linked Census administrative panel data. The sheer breadth and depth of data yield a large number of novel facts and conditional correlations that shed much-needed light on the early reach and economic potential of AI. Here, we distill our findings into a handful of core observations, alongside a discussion of limitations and implications for future work

5.1 Diffusion of AI is low yet concentrated in key segments of the economy

Overall, the view of early AI use among US firms is one of low average diffusion, yet with higher concentration among certain sectors and among a smaller number of very large firms. Exposure to AI among workers is much higher than usage rates among firms: the nation-wide average rate of AI adoption rises from 5.8% to 18.2%, once employment weighting is incorporated.

These rates are much lower than reported in other studies, which lack the representation and coverage of our data set, particularly for the small and medium-sized firms prevalent in many countries, including the US. We do not directly observe impediments to AI adoption, though (lack of) scale clearly plays a role. Indeed, these low averages risk obscuring the fact that most very large firms (over 5,000 employees) report at least some AI use, as well as greater intensity of use.

This pattern of low yet skewed early AI adoption has important implications for its trajectory. Low actual use in production relative to the "hype" likely reflects non-trivial frictions in putting these innovations to work, in practice (e.g., Agrawal et al. 2023; Bresnahan 2023). Yet large firms have disproportionately overcome them, and scale advantages tend to be self-reinforcing. Thus, early disparity in adoption across firm types has the potential to fuel an "AI divide" if such patterns persist. Moreover, uneven firm-level technology adoption has been linked to unequal labor market outcomes (Barth et al. 2023), an area of growing international concern (e.g., Van Reenen 2011). While much current research into the economic implications of AI focuses on labor market outcomes (e.g., at the occupation or skill level), more focus is needed on the role that firms and their competitive dynamics and strategies play in AI's diffusion and longer-term impacts.

Along these lines, our detailed exploration into AI use among startups indicates that early adopters have been concentrated among the types of firms—innovative, with high growth potential—that matter most for economic growth and dynamism. Their rising participation in the economy over time may entail greater presence and impact of AI, overall, potentially challenging existing patterns of concentration or ushering in advantages (e.g., innovation and job creation) that require consideration alongside the risks. Thus our findings, while falling short of sharp predictions, point to key areas of both potential and concern in need of future study.

5.2 "GPT-ness"

Transformative technologies are often identified only with hindsight (Bresnahan and Trajtenberg 1995) and not without controversy (Moser and Nicholas 2004), despite the need to adapt policies and practices to emerging technologies early in their diffusion (Goldfarb et al. 2023). Uninformed AI-related policies and management decisions may be particularly costly (Agrawal et al. 2019), and require evidence-based assessments of AI's reach and potential. While this study does not take a stand on whether AI is a GPT (see Cockburn et al. 2019; Goldfarb et al. 2023; Eloundou et al. 2023), we contribute insights to this discussion. Primarily, we document key characteristics of early AI uptake that meet theoretically necessary conditions for AI to promote reinforcing cycles of diffusion, innovation, and growth across the US economy.

The first of these is breadth of use. Our measure of AI presence captures actual use in production of a variety of AI-related technologies. As such, it distinguishes AI *use* from invention, commercialization, or aspiration (e.g., demand for AI-related skills). Our measure therefore captures a necessary step between technological potential and economic impact that, a) routinely requires costly and often time-consuming process innovation (Bresnahan and Greenstein 1996; McElheran 2015; Feigenbaum and Gross 2021) and, b) is further difficult to observe while it is unfolding. Broad, general use also typically requires uptake within the application-producing sectors of the economy (Bresnahan 2023). Thus use in production in certain sectors is a key leading indicator.

Our first piece of evidence that firms are closing the gap between invention and broad usage is rooted in observed AI use in every sector of the economy, not only those that lead in the supply of AI-related inventions. Higher early adoption within the manufacturing, information, and services sectors—which will supply the goods and services, such as computer software and autonomous robotics, that incorporate AI—is also an important early sign. And, geographic concentration in places that overlap incompletely with "pioneering" tech hubs (Bloom et al. 2021) conveys a similar picture of broadening diffusion. That said, observing change over time based on a consistent set of measures, moving forward, will be needed to refine this interpretation.

The next pattern suggestive of "GPT-ness" is clustering of AI with other important emerging technologies. This has historically been an essential step in technology-driven waves of growth (Rosenberg 1963). Availability of digital information (data) and computing resources (in this case, in the form of cloud computing) appear in our analysis to be central to AI adoption across a very broad swath of firms. High co-occurrence of AI and robotics also points to interdependencies in need of consideration. On the one hand, this suggest that firms are deploying a number of "enabling" technologies in potentially-reinforcing ways. On the other, complementarities among fast-changing digital resources may entail higher costs of adoption and "co-invention" (Bresnahan and Greenstein 1996). Taken together, such mechanisms tend to promote unevenness in the returns to novel technologies, as observed in recent waves of digital transformation (e.g., Brynjolfsson and McElheran 2016; Brynjolfsson et al. 2021). We unfortunately do not directly observe the costs of adopting any of these technologies. A useful priority for future work would be to understand heterogeneity in both adoption costs and performance associated with AI and related digital technologies.

The last key indicator of "GPT-ness" is the observed conditional correlation between AI use and a direct measure of recent process innovation at the firm. In different samples and controlling for a large number of correlates of AI use (not to mention fine-grained industry controls), firms that innovated within the prior three years in a range of business processes were significantly more likely to use AI. This is again consistent with frictions in AI adoption at the process level, yet with evidence that they are nevertheless being overcome, in practice. We have limited insight, however, into what dimensions of process-technology alignment are most essential and/or challenging. Our most-detailed insights are further most-robust for younger firms, leaving open questions around the adaptation activities of much larger and older incumbent firms in need of study.

5.3 Firms associated with "high-growth" entrepreneurship adopt AI

Our interest in firm dynamics led us to examine AI use more closely in a large sample of startup firms (75,000) for which the ABS provides unprecedented visibility to details of the organizational context. Among startups, early indicators of high growth potential are not only empirically linked to actual revenue growth in the administrative data, they are also robustly correlated with AI use. Factors associated with the small number of highest-growth "gazelles," such as VC funding, owning patents, and high startup capitalization, have the largest marginal association (all else equal) with the likelihood of adoption. VC-funded startups tend to be high-growth businesses even before VC-funding, and many venture capitalists have the skills and expertise to assess the underlying promise of a startup. Hence, VC funding could be a signal of firm potential not captured by other observables in our analysis. However, we also do not rule out reverse causality, whereby VC influence promotes AI adoption. Either way, the characteristics of VC-funded firms appear highly complementary to AI use in ways that have been heretofore unobserved.

As part of our interest in young high-potential firms, we explored the role of firm leaders and their characteristics in AI adoption. This revealed that more educated, more experienced, and younger owners (typically founders) are more likely to preside over AI adoption at their firms. Our finding that firms with owners younger than 35 are more likely to adopt AI is consistent with evidence that younger entrepreneurs tend to be more sophisticated users of advanced technology and may be more open to adopting these technologies in their businesses. Our estimates, however, flag an age range that is younger than the mean age of high-growth founders (Azoulay et al. 2020) or beneficiaries of digitalization (Barth et al. 2023) found in prior work.

Serial entrepreneurship, which has been primarily studied in relationship to firm performance, conditionally predicts AI use. This suggests a potential mechanism underlying uneven performance gains in the digital age beyond the scope of our study, yet worthy of deeper understanding.

Our finding that the motivations of owners and founders are correlated with AI use is novel to the literature on technology diffusion and its implications. We document that high-growth and innovative firms are not only led by individuals motivated by bringing new ideas to market but also by those who report prosocial values such as helping or becoming involved in the community. In turn, such motivations are highly correlated with AI use. Other motivations for founding or owning firms matter, but there is a clear distinction between markers of growth-oriented and so-called "lifestyle" entrepreneurship when it comes to AI use. This again suggests an outsized impact of AI compared to its low early average adoption, while pointing to potential guardrails regarding its specific applications, in practice. A useful direction for future work would be to explore the link between owner motivations and "ethical AI" (e.g., Balagopalan et al. 2023) in practice.

5.4 Firm strategies matter for AI use

Innovation and growth-oriented business strategies, also difficult to observe at scale, significantly predict both early revenue growth and AI use. While process innovation dominates product innovation as a predictor of AI use, product innovation and reliance on formal IP are key correlates of startup use of AI (despite their high baseline prevalence and conditioning on a large number of other early-growth indicators). Growth-oriented strategy has a smaller but nevertheless significant correlation with AI use, even conditioning on actual early revenue growth.

Indeed, conditioning directly on a large number of markers of high-growth potential reveals a robust relationship between AI use and growth *per se*, among startups. A possible interpretation is that firms with a demonstrated upward trajectory, early on, may have a greater demand—as well as the requisite resources and organizational complements– to exploit AI. Even so, our analysis of startups, as well as of firms for which we can assess revenue growth later in life, points to an important relationship between growth and adoption that is consistent with performance gains due to AI use. Recall that while we measure AI use in the cross-section, we have a panel of administrative performance data to deepen insights, here. That said, establishing causality is beyond the scope of our study, and we hope that these findings may provide the foundation for follow-on efforts to understand causality between AI use and firm performance.

5.5 Geographic disparity in the adoption of AI is pronounced

Our fifth main stylized fact is substantial geographic concentration in the adoption and use of AI. Consistent with our in-depth exploration of young firms, above, our findings emphasize key "hubs" of AI use in production among single-unit startups. We find the rate of AI use to be higher in CBSAs located in the southern and western parts of the US. Concentration in known technology hubs (Kerr and Robert-Nicoud 2020; Bloom et al. 2021), such as Silicon Valley and the Research Triangle, is consistent with clustering of job postings relating to AI and the imporance of proximity to academic centers (e.g., Muro and Liu 2021; Bloom et al. 2021; Bessen et al. 2021). Certain large metro areas such as Nashville, San Antonio, Las Vegas, New Orleans, San Diego, and Tampa, also show high concentrations of AI use. The concentration of AI in these already-leading locations suggests an "AI divide" that may be further reinforced if new AI-using, high-growth startups continue to be attracted to locations with already-high AI activity.

That said, our ranking of AI clusters overlaps incompletely with prior work. And, expanding our lens to include a large number of older single-unit firms,[29] we see clustering of AI use in a broader set of locations. While uncovering specific drivers of early AI geography is beyond the scope of our study, our interpretation is that the forces promoting agglomeration in *use* may be distinct from those promoting co-location for invention or early commercialization, as identified in measures based on patents, publications, and/or demand for AI-related skills. Further, when we weight by employment, a number of less-obvious locales attain "hub status," primarily in the Midwest and Mid-Atlantic. These include Louisville, Columbus, Austin, and Atlanta. Discussions around the "future of work" as AI diffuses may need to address different considerations for firms seeking to develop or supply AI-related technologies, versus those whose activities and sheer size expose a broader range of employees to AI at work.

5.6 Characterizing AI adoption poses a significant measurement challenge

While not one of our five main stylized facts, a consideration that runs throughout our study is the recognition that typically-observable firm characteristics leave unexplained a large fraction of the variation in AI adoption across firms. Even with high-dimensional controls (state–industry controls at the 6-digit NAICS level), along with controls for firm age and owner gender, we can only explain around 21% of the variation in adoption (see Table 7). This low explanatory power despite detailed covariates indicates that significant unobserved heterogeneity remains in AI adoption among US firms. Including additional, harder-to-measure factors from the ABS related to leadership, capitalization, and innovative strategies contributes significantly to the explanatory power (an increase of about 10% from the baseline specification). Nevertheless, a large amount of variation is left unexplained, underscoring the importance of idiosyncratic firm-level unobserved factors, such as the specific product or service being offered by a firm or the types of processes and tasks used in production. This points to a number of important areas for follow-on data collection and research, as well as the recognition that choices taken around measurement may have significant implications for inference.

6 Conclusion

As AI advances and becomes more integrated into the workplace, there is a growing debate about whether it will increase productivity, what the effects on the workforce will be, or even whether AI will spark a revolution in business processes and models. Central to these debates is obtaining reliable measures of AI use across the U.S. firm population, identifying relevant adopter characteristics, and assessing worker exposure to AI.

The 2018 ABS technology module addresses this data need along a number of dimensions, providing detailed information about the early diffusion of AI and related technologies across a nationally representative set of firms, further linked to detailed information on industry, age, and revenue over the firm life cycle. It further provides unprecedented insights into the technological and organizational context of AI use, particularly among high-potential US startups.

As touched on throughout, our study is not without limitations. While our "headline" adoption statistics are robust and unusually representative of the entire US economy, our most-nuanced characterization of AI users leans on a particular subsample of young firms. This sample is large and covers important swaths of the firm size and industry distributions, yet it is also subject to non-trivial ownership and age restrictions. Ultimately, our approach trades off unusual visibility to early-stage entrepreneurial activity—as well as changes across the life cycle—for insights that may be less-applicable to the publicly traded firms predominant in related research. To the extent that we seek to signpost AI's future trajectory, in addition to documenting its early baseline, this tradeoff is worth making. Our findings are also conditional on survival: any firms that adopted AI and failed before 2018 are missing from our data.[30]

The exploration of AI's leading edge provided here arguably raise more questions than they answer. Among the many patterns we document, the potential for an "AI divide" between different types of firms and geographies is visible—and in need of careful assessment. Pinning down causal relationships is also beyond the reach of this first cross-section of early data. Our hope is that future collections can build a panel of stable measures and identify exogenous variation in AI use for this important endeavor. That said, the patterns we uncover concerning the who, what, and where of early adoption of AI may support a more evidence-based response to current technological advances and economic challenges and identify promising paths for future work.

Endnotes

1 For reference, the quinquennial Economic Census typically has a response rate of approximately 80% for single-unit firms (Dinlersoz and Klimek 2011). The Community Innovation Survey in Europe reported response rates of 25–35% in Germany, resulting in approximately 5.5 thousand firms (Czarnitzki et al. 2023). The Eurostat survey on AI had response rates of 5–19% across 27 EU countries, yielding over 9.6 thousand firms (Hoffreumon et al. 2023).

2 The AI-related technologies covered are: automated-guided vehicles, machine learning, machine vision, natural language processing, and voice recognition. See Table 2. To best capture the extensive margin of AI use across industries, we combine usage rates across AI-related technologies. That said, our findings largely represent variation in the use of "machine learning" in production.

3 Young firms are disproportionately single-unit, and hence straightforward to situate geographically.

4 Studies of AI use and impacts outside of the US have been on the rise, including in Europe (Czarnitzki et al. 2023; Hoffreumon et al. 2023), China (Lu et al. 2023; Beraja et al. 2023), Canada (Alexopoulos and Cohen 2018) and in the context of international trade (Goldfarb and Trefler 2018; Brynjolfsson et al. 2019).

5 See, for example Galdon-Sanchez et al. (2022) on the value of "big data" for small bricks-and-mortar firms in Spain.

6 E.g., Catalini et al. (2019); Akcigit et al. (2022); Lerner and Nanda (2020); Botelho et al. (2021).

7 See Brynjolfsson and Milgrom (2013) for more on complementarity theory and a review of the literature.

8 Such details are captured via responses from "primary owners" to the ABS, which excludes publicly traded firms, estates, trusts, government and tribal entities, associations, membership clubs, cooperatives and foreign entities.

9 Specifically, we recalculate the sample weights by stratifying the firms in the 2017 LBD and our final sample of firms in the ABS on firm size, age, and industry. These strata are defined by the 19 two-digit NAICS sectors and the 12 firm size and 12 firm age groups used in the BDS. We weight according to these more standard firm characteristics because of the unique sampling frames of the ABS (see Section B.1 in our data appendix).

10 Distributions taken from Zolas et al. (2020) show that nearly 70 percent of ABS firms, and more than 3/4 once weighted, have fewer than 10 employees, while the age distribution is more uniform.

11 Other technologies queried include augmented reality, robotics, touchscreens, automated storage and retrieval systems, and radio frequency identification (RFID). While some of these technologies may also contain an AI-related component, they are less likely to rely heavily on AI compared to the ones we consider to be AI-intensive. For instance, the 2018 ABS definition of robotics (a technology that has increasingly relied on AI in recent years) does not mention autonomy, but rather only reprogrammability (i.e., "traditional" industrial robots).

12 Also, the rates of missing data vary by specific technology, which would affect the denominator in calculating rates.

13 The rise in usage rates with firm size could be partly a "mechanical" function of scale. If we conceptualize larger firms as random collections of smaller ones and extrapolate usage rates as a function of number of employees, we would expect even higher usage rates among larger firms.

14 E.g., this controls for the distinction between manufacturing hardwood versus softwood veneer (NAICS 321211 vs. 321212) and bicycles vs. armored vehicles and tanks (336991 vs. 336992); newspaper publishers vs. greeting card publishers (513110 vs. 514191); and casino hotels versus bed-and-breakfast inns (721120 vs. 721191).

15 Since the LBD only extends back to 1976, the age distribution in our samples is truncated from above at 42 years, with all firms of age 42+ being assigned to the highest age percentile group. As a result, coefficient for the highest age percentile should be interpreted with caution.

16 We see similar patterns of increasing overlap as firms get larger. Figure A1 plots the share of AI and robotics adopters across the different size categories. Among AI adopters, the share of firms also adopting robotics rises with size. However, somewhat interestingly, we find the opposite pattern among robotics adopters. The share of robotics users that also use AI declines from 65.6% of robotics adopters in the smallest size category to less than half (49.3%) of firms with 20 or more employees. These patterns suggest that any underlying complementarities between AI and other technologies may have complex interactions with firm characteristics such as scale.

17 For details on the owner and firm characteristics and how they are collected in the ABS, see Data Appendix B, as well as (Zolas et al. 2020). This subsample is restricted to firms with information on the primary owner, thus necessarily excluding any publicly-owned firms. For most analyses involving these measures, we further require revenue information. Column 3 of Table 1 describes the relevant subsample, which is contingent on a successful match to the LBD and its introduction of revenue data in 1997. This excludes firms that were older than 20 in our reference year of 2017. These limitations notwithstanding, the coverage and size of the data set (209,000 firms), along with its focus on the use of AI in production, are significant in this literature. See Section 2 for details and references.

18 For instance, in row 6, only 24.7% of such owners were 55 years or older (column 2), compared to 26.6% of the overall startup sample (column 1). That said, roughly 55% of all startups (AI using or otherwise) fell in the 35-54 range bracketing "peak entrepreneurship" (Azoulay et al. 2020).

19 CBSAs are geographic areas consisting of one or more counties (or equivalents) with an urban center of at least 10,000 people plus adjacent counties that are socioeconomically tied to the urban center by commuting. For multi-unit firms in this "startup subsample," we assign a single CBSA by taking the maximum employment within a firm-zip code.

20 Because a sample of single-unit firms differs somewhat from our nationally representative baseline sample, we modify our firm weights. We do not, however, directly incorporate geographic information when generating our weights. For comments on incorporating geography into tabulation weights, see Section B.4 in our data appendix.

21 Because we only require location information, this sample includes startups that do not appear in our core startup sample from column 4, Table 1. However, because of the granularity of CBSAs and to protect confidentiality, we only report information for the 30,000 startups in the largest CBSAs (those with a population of 1 million or more).

22 Because Durham–Chapel Hill does not meet our population requirement to be designated as a "large" CBSA, this piece of the Research Triangle is absent from Figure 5a and Table A2. However, Figure 6a shows that Durham–Chapel Hill has a relatively high usage rate of AI when considering all single-unit firms.

23 As shown in Figure 6a, when considering all single-unit firms there are also pockets of high adoption in some smaller CBSAs such as Provo, UT and Eugene, OR—including Madison, WI and Akron, OH when weighting by employment (see Figure 6b).

24 This analysis restricts on firms that survived through 2017 and responded to the survey in 2018: differential survival rates of adopters versus non-adopters prior to 2017 would introduce biases whose sign is difficult to predict. That said, most studies of technology adoption and the impact of impact technology on performance are vulnerable to this concern (see Jin and McElheran 2017 for a discussion and evidence).

25 "Unable to find employment" is also negative, statistically significant, and relatively large in magnitude. However it has little explanatory power in the AI usage regressions, to follow. A similar lack of explanatory power shows up for "Wanted to be own boss," so we abstract away from these in our core analyses.

26 Recall that we do not claim a causal relationship. In fact, faster-growing startups may be more able to use and benefit from AI if the initial costs of adoption or the scale of training data are large.

27 As described in Section 3.6.1, we find (unconditionally) that the use of AI-related technologies is negatively correlated with "Flexible hours" and "Balance work and family" reasons, and it is positively correlated with "Best avenue for ideas/goods/services" and "Help and/or become more involved in community" (Table A3). For simplicity we thus retain our "lifestyle" reason indicator in Table 7.

28 This "set" of covariates, examined on their own analogous to columns 2–4, yield nearly identical effects and hence are combined, here, for brevity.

29 Conditional on data availability, per Table 1 and Appendix B.

30 This limitation is common in digital technology research and underscores the importance of collecting data on new technologies as early as possible (e.g., McElheran 2018).

References

Acemoglu, D., U. Akcigit, H. Alp, N. Bloom, and W. Kerr (2018). Innovation, reallocation, and growth. *American Economic Review 108*(11), 3451–3491.

Acemoglu, D., D. Autor, J. Hazell, and P. Restrepo (2022). Artificial intelligence and jobs: evidence from online vacancies. *Journal of Labor Economics 40*(S1), S293–S340.

Acemoglu, D. and P. Restrepo (2020). The wrong kind of ai? artificial intelligence and the future of labour demand. *Cambridge Journal of Regions, Economy and Society 13*(1), 25–35.

Agrawal, A., J. Gans, and A. Goldfarb (Eds.) (2019). *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press.

Agrawal, A., J. S. Gans, and A. Goldfarb (2023). Artificial intelligence adoption and system-wide change. *Journal of Economics & Management Strategy*.

Akcigit, U., E. Dinlersoz, J. Greenwood, and V. Penciakova (2022). Synergizing Ventures. *Journal of Economic Dynamics and Control, (forthcoming)*.

Alcácer, J. and W. Chung (2014). Location strategies for agglomeration economies. *Strategic Management Journal 35*(12), 1749–1761.

Alekseeva, L., J. Azar, M. Giné, S. Samila, and B. Taska (2021). The demand for AI skills in the labor market. *Labour Economics 71*(C).

Alexopoulos, M. and J. Cohen (2018). Canadian productivity growth, secular stagnation, and technological change. *International Productivity Monitor* (35), 113–137.

Aral, S., E. Brynjolfsson, and L. Wu (2012). Three-way complementarities: Performance pay, human resource analytics, and information technology. *Management Science 58*(5), 913–931.

Armbrust, M., A. Fox, R. Griffith, A. D. Joseph, R. Katz, A. Konwinski, G. Lee, D. Patterson, A. Rabkin, I. Stoica, and M. Zaharia (2010, apr). A view of cloud computing. *Commun. ACM 53*(4), 50–58.

Autor, D., D. Dorn, L. F. Katz, C. Patterson, and J. V. Reenen (2020). The Fall of the Labor Share and the Rise of Superstar Firms ["Automation and New Tasks: How Technology Displaces and Reinstates Labor"]. *The Quarterly Journal of Economics 135*(2), 645–709.

Azoulay, P., B. F. Jones, J. D. Kim, and J. Miranda (2020). Age and High-Growth Entrepreneurship. *American Economic Review: Insights 2*(1), 65–82.

Babina, T., A. Fedyk, A. X. He, and J. Hodson (2022). Firm investments in artificial intelligence technologies and changes in workforce composition. *Available at SSRN 4060233*.

Balagopalan, A., D. Madras, D. H. Yang, D. Hadfield-Menell, G. K. Hadfield, and M. Ghassemi (2023). Judging facts, judging norms: Training machine learning models to judge humans requires a modified approach to labeling data. *Science Advances 9*(19), eabq0701.

Bartel, A., C. Ichniowski, and K. Shaw (2007). How does information technology affect productivity? plant-level comparisons of product innovation, process improvement, and worker skills. *The quarterly journal of Economics 122*(4), 1721–1758.

Barth, E., J. C. Davis, R. B. Freeman, and K. McElheran (2023). Twisting the demand curve: Digitalization and the older workforce. *Journal of Econometrics 233*(2), 443–467.

Beraja, M., A. Kao, D. Y. Yang, and N. Yuchtman (2023). Ai-tocracy. *The Quarterly Journal of Economics 138*(3), 1349–1402.

Bessen, J., I. Cockburn, and J. Hunt (2021). Is distance from innovation a barrier to the adoption of artificial intelligence? *Boston University Working Paper*.

Bessen, J., S. M. Impink, L. Reichensperger, and R. Seamans (2022). The role of data for ai startup growth. *Research Policy 51*(5), 104513.

Black, S. E. and L. M. Lynch (2001). How to compete: the impact of workplace practices and information technology on productivity. *Review of Economics and statistics 83*(3), 434–445.

Bloom, N., T. A. Hassan, A. Kalyani, J. Lerner, and A. Tahoun (2021). The diffusion of disruptive technologies. Technical report, National Bureau of Economic Research.

Bloom, N., R. Sadun, and J. Van Reenen (2012). Americans do it better: Us multinationals and the productivity miracle. *American Economic Review 102*(1), 167–201.

Botelho, T. L., D. Fehder, and Y. Hochberg (2021). Innovation-Driven Entrepreneurship. NBER Working Papers 28990, National Bureau of Economic Research, Inc.

Bresnahan, T. (2023). What innovation paths for ai to become a gpt? *Journal of Economics & Management Strategy*.

Bresnahan, T. and S. Greenstein (1996). Technical progress and co-invention in computing and in the uses of computers. *Brookings Papers on Economic Activity. Microeconomics 1996*, 1–83.

Bresnahan, T. F., E. Brynjolfsson, and L. M. Hitt (2002). Information technology, workplace organization, and the demand for skilled labor: Firm-level evidence. *The Quarterly Journal of Economics 117*(1), 339–376.

Bresnahan, T. F. and M. Trajtenberg (1995). General purpose technologies 'engines of growth'? *Journal of econometrics 65*(1), 83–108.

Brynjolfsson, E. (2022). The turing trap: The promise & peril of human-like artificial intelligence. *Daedalus 151*(2), 272–287.

Brynjolfsson, E. and L. Hitt (1996). Paradox lost? firm-level evidence on the returns to information systems spending. *Management science 42*(4), 541–558.

Brynjolfsson, E. and L. M. Hitt (2000). Beyond computation: Information technology, organizational transformation and business performance. *Journal of Economic perspectives 14*(4), 23–48.

Brynjolfsson, E., X. Hui, and M. Liu (2019). Does machine translation affect international trade? evidence from a large digital platform. *Management Science 65*(12), 5449–5460.

Brynjolfsson, E., W. Jin, and K. McElheran (2021). The power of prediction: predictive analytics, workplace complements, and business performance. *Business Economics 56*(4), 217–239.

Brynjolfsson, E. and K. McElheran (2016). The rapid adoption of data-driven decision-making. *American Economic Review 106*(5), 133–39.

Brynjolfsson, E. and K. McElheran (2019). Data in action: data-driven decision making and predictive analytics in us manufacturing. *Rotman School of Management Working Paper* (3422397).

Brynjolfsson, E. and P. Milgrom (2013). Complementarity in organizations. *The handbook of organizational economics*, 11–55.

Brynjolfsson, E., T. Mitchell, and D. Rock (2018). What Can Machines Learn, and What Does It Mean for Occupations and the Economy? *AEA Papers and Proceedings 108*, 43–47.

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

Camuffo, A., A. Gambardella, A. Pignataro, et al. (2022). Framing strategic decisions in the digital world. *Strategic Management Review, forthcoming*.

Cassar, L. and S. Meier (2018). Nonmonetary incentives and the implications of work as a source of meaning. *Journal of Economic Perspectives 32*(3), 215–38.

Catalini, C., J. Guzman, and S. Stern (2019). Hidden in Plain Sight: Venture Growth with or without Venture Capital. NBER Working Papers 26521, National Bureau of Economic Research, Inc.

Chari, V. V. and H. Hopenhayn (1991). Vintage human capital, growth, and the diffusion of new technology. *Journal of political Economy 99*(6), 1142–1165.

Chava, S., A. Oettl, A. Subramanian, and K. V. Subramanian (2013). Banking deregulation and innovation. *Journal of Financial economics 109*(3), 759–774.

Chow, M. C., T. C. Fort, C. Goetz, N. Goldschlag, J. Lawrence, E. R. Perlman, M. Stinson, and T. K. White (2021). Redesigning the longitudinal business database. Working Paper 28839, National Bureau of Economic Research.

Chui, M. and S. Malhotra (2018). Ai adoption advances, but foundational barriers remain. *Mckinsey and Company*.

Cockburn, I. M., R. Henderson, and S. Stern (2019). The impact of artificial intelligence on innovation. *The economics of artificial intelligence: An agenda*, 115–152.

Cohen, W. M. and S. Klepper (1996). Firm size and the nature of innovation within industries: the case of process and product r&d. *The review of Economics and Statistics*, 232–243.

Cowgill, B. and C. E. Tucker (2019). Economics, fairness and algorithmic bias. *preparation for: Journal of Economic Perspectives*.

Criscuolo, C., P. N. Gal, and C. Menon (2014). The dynamics of employment growth: New evidence from 18 countries.

Czarnitzki, D., G. P. Fernández, and C. Rammer (2023). Artificial intelligence and firm-level productivity. *Journal of Economic Behavior and Organization 211*, 188–205.

David, P. A. (1990). The dynamo and the computer: an historical perspective on the modern productivity paradox. *The American Economic Review 80*(2), 355–361.

Decker, R. A., J. Haltiwanger, R. S. Jarmin, and J. Miranda (2017). Declining Dynamism, Allocative Efficiency, and the Productivity Slowdown. *American Economic Review 107*(5), 322–326.

Decker, R. A., J. Haltiwanger, R. S. Jarmin, and J. Miranda (2020). Changing Business Dynamism and Productivity: Shocks versus Responsiveness. *American Economic Review 110*(12), 3952–3990.

Delgado, M., M. E. Porter, and S. Stern (2010). Clusters and Entrepreneurship. *Journal of Economic Geography 10*(4), 495–518.

Deloitte (2018). State of ai in the enterprise.

Dinlersoz, E., N. Goldschlag, A. Myers, and N. Zolas (2021). An Anatomy of US Firms Seeking Trademark Registration. In *Measuring and Accounting for Innovation in the Twenty-First Century*, NBER Chapters, pp. 183–228. National Bureau of Economic Research, Inc.

Dinlersoz, E., N. Goldschlag, M. Yorukoglu, and N. Zolas (2023). On the Role of Trademarks: From Micro Evidence to Macro Outcomes. Center for Economic Studies Working Paper 23-16, U.S. Census Bureau.

Dinlersoz, E. and S. Klimek (2011). Modeling Single Establishment Firm Returns to the 2007 Economic Census. Working Paper CES-WP-11-28, U.S. Census Bureau.

Dinlersoz, E. and Z. Wolf (2018). Automation, Labor Share, and Productivity: Plant-Level Evidence from U.S. Manufacturing. Working Papers 18-39, Center for Economic Studies, U.S. Census Bureau.

Dixon, J., B. Hong, and L. Wu (2021). The robot revolution: Managerial and employment consequences for firms. *Management Science 67*(9), 5586–5605.

Eloundou, T., S. Manning, P. Mishkin, and D. Rock (2023). Gpts are gpts: An early look at the labor market impact potential of large language models. *arXiv preprint arXiv:2303.10130*.

Feigenbaum, J. J. and D. P. Gross (2021). Organizational frictions and increasing returns to automation: Lessons from at&t in the twentieth century. Technical report, National Bureau of Economic Research.

Felten, E., M. Raj, and R. Seamans (2021). Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. *Strategic Management Journal 42*(12), 2195–2217.

Forman, C., A. Goldfarb, and S. Greenstein (2005). How did location affect adoption of the commercial internet? global village vs. urban leadership. *Journal of urban Economics 58*(3), 389–420.

Forman, C., A. Goldfarb, and S. Greenstein (2012). The internet and local wages: A puzzle. *American Economic Review 102*(1), 556–75.

Forman, C., A. Goldfarb, and S. Greenstein (2016). Agglomeration of invention in the bay area: Not just ict. *American Economic Review 106*(5), 146–51.

Galdon-Sanchez, J. E., R. Gil, and G. Uriz Uharte (2022). The value of information in competitive markets: The impact of big data on small and medium enterprises. *Available at SSRN 4022675*.

Gambardella, A., S. Heaton, E. Novelli, and D. J. Teece (2021). Profiting from enabling technologies? *Strategy Science 6*(1), 75–90.

Ganguli, I., M. Huysentruyt, and C. Le Coq (2021). How do nascent social entrepreneurs respond to rewards? a field experiment on motivations in a grant competition. *Management Science 67*(10), 6294–6316.

Gans, J. S., S. Stern, and J. Wu (2019). Foundations of entrepreneurial strategy. *Strategic Management Journal 40*(5), 736–756.

Ghemawat, P. (1991). *Commitment*. Simon and Schuster.

Goldfarb, A., B. Taska, and F. Teodoridis (2023). Could machine learning be a general purpose technology? a comparison of emerging technologies using data from online job postings. *Research Policy 52*(1), 104653.

Goldfarb, A. and D. Trefler (2018). Artificial intelligence and international trade. In *The economics of artificial intelligence: an agenda*, pp. 463–492. University of Chicago Press.

Griliches, Z. (1957). Hybrid corn: An exploration in the economics of technological change. *Econometrica, Journal of the Econometric Society*, 501–522.

Guzman, J. (2019). *Go west young firm: agglomeration and embeddedness in startup migrations to Silicon Valley*. SSRN.

Guzman, J., J. J. Oh, and A. Sen (2020). What motivates innovative entrepreneurs? evidence from a global field experiment. *Management Science 66*(10), 4808–4819.

Guzman, J. and S. Stern (2015). Where is silicon valley? *Science 347*(6222), 606–609.

Guzman, J. and S. Stern (2016). Nowcasting and placecasting entrepreneurial quality and performance. In *Measuring entrepreneurial businesses: Current knowledge and challenges*, pp. 63–109. University of Chicago Press.

Guzman, J. and S. Stern (2020). The State of American Entrepreneurship: New Estimates of the Quantity and Quality of Entrepreneurship for 32 US States, 1988–2014. *American Economic Journal: Economic Policy 12*(4), 212–243.

Haltiwanger, J., R. S. Jarmin, R. Kulick, and J. Miranda (2016). High Growth Young Firms: Contribution to Job, Output, and Productivity Growth. In *Measuring Entrepreneurial Businesses: Current Knowledge and Challenges*, NBER Chapters, pp. 11–62. National Bureau of Economic Research, Inc.

Haltiwanger, J., R. S. Jarmin, and J. Miranda (2013). Who creates jobs? small versus large versus young. *Review of Economics and Statistics 95*(2), 347–361.

Hoffreumon, C., C. Forman, and N. van Zeebroeck (2023). Make or buy your artificial intelligence? testing for complementarities in technology sourcing. *Cornell University Working Paper*.

Hong, S., D. H. Hsu, and V. Subramanian (2023). The evolution of u.s. entrepreneurial and innovation ecosystems in artificial intelligence. *Wharton School Working Paper*.

Hurst, E. and B. W. Pugsley (2011). What Do Small Businesses Do? *Brookings Papers on Economic Activity 42*(2 (Fall)), 73–142.

Iansiti, M. and K. R. Lakhani (2020). *Competing in the age of AI: Strategy and leadership when algorithms and networks run the world*. Harvard Business Press.

Ichniowski, C., K. L. Shaw, and G. Prennushi (1995). The effects of human resource management practices on productivity.

Jaffe, A. B., M. Trajtenberg, and R. Henderson (1993). Geographic localization of knowledge spillovers as evidenced by patent citations. *The Quarterly Journal of Economics 108*(3), 577–598.

Jin, W. and K. McElheran (2017). Economies before scale: survival and performance of young plants in the age of cloud computing. *Rotman School of Management Working Paper* (3112901).

Kapoor, R. and D. J. Teece (2021). Three faces of technology's value creation: Emerging, enabling, embedding. *Strategy Science 6*(1), 1–4.

Kazakova, S. S., A. A. Dunne, D. D. Bijwaard, J. Gossé, C. Hoffreumon, N. van Zeebroeck, et al. (2020). European enterprise survey on the use of technologies based on artificial intelligence. Technical report, ULB–Universite Libre de Bruxelles.

Kerr, W. R. and F. Robert-Nicoud (2020). Tech clusters. *Journal of Economic Perspectives 34*(3), 50–76.

Kueng, L., M.-J. Yang, and B. Hong (2014). Sources of firm life-cycle dynamics: differentiating size vs. age effects. Technical report, National Bureau of Economic Research.

Lafontaine, F. and K. Shaw (2016). Serial Entrepreneurship: Learning by Doing? *Journal of Labor Economics 34*(S2), 217–254.

Lerner, J. and R. Nanda (2020). Venture Capital's Role in Financing Innovation: What We Know and How Much We Still Need to Learn. *Journal of Economic Perspectives 34*(3), 237–261.

Levine, R. and Y. Rubinstein (2018). Selection into entrepreneurship and self-employment. Technical report, National Bureau of Economic Research.

Li, J. F. (2023). Jump starting the ai engine: The complementary role of data and management practices. *Available at SSRN 4495624*.

Lu, Y., G. M. Phillips, and J. Yang (2023). The impact of cloud computing and ai on industry dynamics and competition. *Available at SSRN 4480570*.

McElheran, K. (2015). Do market leaders lead in business process innovation? the case (s) of e-business adoption. *Management science 61*(6), 1197–1216.

McElheran, K. (2018). Economic measurement of ai. Technical report, National Bureau of Economic Research.

McElheran, K. and W. Jin (2023). The strategic fit of it resources in the age of cloud computing. *University of Toronto Working paper*.

McElheran, K., S. Ohlmacher, and M.-J. Yang (2019). Strategy and Structured Management. Working paper, University of Toronto.

Miric, M., N. Jia, and K. G. Huang (2023). Using supervised machine learning for large-scale classification in management research: The case for identifying artificial intelligence patents. *Strategic Management Journal 44*(2), 491–519.

Moser, P. and T. Nicholas (2004). Was electricity a general purpose technology? evidence from historical patent citations. *American Economic Review 94*(2), 388–394.

Muro, M. and S. Liu (2021). The geography of ai, which cities will drive the artificial intelligence revolution? *Brookings Institution*.

Petralia, S. (2020). Mapping general purpose technologies with patent data. *Research Policy 49*(7), 104013.

Porter, M. E. (1980). Techniques for analyzing industries and competitors. *Competitive Strategy. New York: Free*.

PwC (2019). 2019 ai predictions.

Raj, M. and R. Seamans (2019). Primer on artificial intelligence and robotics. *Journal of Organization Design 8*, 1–14.

Rathje, J. M. and R. Katila (2021). Enabling technologies and the role of private firms: A machine learning matching analysis. *Strategy Science 6*(1), 5–21.

Rosenberg, N. (1963). Technological change in the machine tool industry, 1840–1910. *The journal of economic history 23*(4), 414–443.

Samila, S. and O. Sorenson (2010, December). Venture capital as a catalyst to commercialization. *Research Policy 39*(10), 1348–1360.

Shah, S. K., R. Agarwal, and R. Echambadi (2019). Jewels in the crown: Exploring the motivations and team building processes of employee entrepreneurs. *Strategic Management Journal 40*(9), 1417–1452.

Spulber, D. F. (2011a). How entrepreneurs affect the rate and direction of inventive activity. In *The Rate and Direction of Inventive Activity Revisited*, pp. 277–315. University of Chicago Press.

Spulber, D. F. (2011b). Should business method inventions be patentable? *Journal of Legal Analysis 3*(1), 265–340.

Tambe, P. (2014). Big data investment, skills, and firm value. *Management science 60*(6), 1452–1469.

Tambe, P., L. Hitt, D. Rock, and E. Brynjolfsson (2020). Digital capital and superstar firms. Technical report, National Bureau of Economic Research.

Tambe, P., L. M. Hitt, and E. Brynjolfsson (2012). The extroverted firm: How external information practices affect innovation and productivity. *Management Science 58*(5), 843–859.

Trajtenberg, M. (2018). Ai as the next gpt: a political-economy perspective. Technical report, National Bureau of Economic Research.

Van Reenen, J. (2011). Wage inequality, technology and trade: 21st century evidence. *Labour economics 18*(6), 730–741.

Webb, M. (2019). The impact of artificial intelligence on the labor market. *Available at SSRN 3482150*.

Wu, L., L. Hitt, and B. Lou (2020). Data analytics, innovation, and firm productivity. *Management Science 66*(5), 2017–2039.

Wu, L., B. Lou, and L. Hitt (2019). Data analytics supports decentralized innovation. *Management Science 65*(10), 4863–4877.

Zolas, N., Z. Kroff, E. Brynjolfsson, K. McElheran, D. N. Beede, C. Buffington, N. Goldschlag, L. Foster, and E. Dinlersoz (2020). Advanced Technologies Adoption and Use by U.S. Firms: Evidence from the Annual Business Survey. NBER Working Papers 28290, National Bureau of Economic Research, Inc.

Tables

Table 1: Summary Statistics of Firm Size and Age by Sample

	(1) ABS–LBD Linked Sample	(2) Baseline Sample	(3) Owner & Revenue Sample	(4) Startup Sample	(5) Startup Sample in Spec. Ind.
Employment (Unweighted)	89.32	61.69	17.68	10.39	9.11
Employment (Weighted)	26.28	20.77	9.78	7.4	6.91
Age (Unweighted)	16.33	16.60	8.94	2.87	2.87
Age (Weighted)	15.61	15.94	8.66	2.82	2.83
Observations (Unweighted)	573,000	447,000	209,000	75,000	28,000
Observations (Weighted)	5,180,000	4,050,000	1,970,000	740,000	228,000

Notes: Tabulations based on 2018 ABS data linked with the 2017 Longitudinal Business Database (LBD) data for size and age. Firms that did not respond to any of the 2018 ABS survey are excluded. See Zolas et al. (2020) for further details and breakdown of the size and age distribution for the full sample of linked ABS–LBD firms. Weights are computed by stratifying the firms in the 2017 LBD and our final sample of firms in the ABS on firm size, age, and industry. These strata are defined by the 19 two-digit NAICS sectors and the 12 firm size and 12 firm age groups used in the BDS. Samples in columns 2–4 are restricted by needing to contain information on the primary owner and are born on or after 1997, the first year that the LBD contains firm-level revenue measures. "Startups" are defined as being born on or after 2012. Specific Industries refers to a subset of 2-digit NAICS sectors that are more engaged in AI usage: Manufacturing ("31–33"), Information ("51"), Professional and Scientific Services ("54") and Healthcare ("62").

Table 2: Definitions of Technologies

Augmented reality	Technology that provides a view of a real-world environment with computer-generated overlays.
Automated guided vehicles (AGV) or AGV systems	A computer-controlled transport vehicle that operates without a human driver. AGVs navigate facilities through the use of software and sensors.
Automated storage and retrieval systems	Technology that locates, retrieves, and replaces items from predetermined storage locations.
Machine learning	Computer algorithms that use data to improve their predictive performance without being reprogrammed.
Machine vision	Technology used to provide image-based automatic inspection, recognition or analysis.
Natural language processing	Technology that allows a computer to process human speech or text.
Radio-frequency identification (RFID) system	A system of tags and readers used for identification and tracking. Tags store information and transmit them using radio waves. Readers maybe be mobile or fixed in place.
Robotics	Reprogrammable machines capable of automatically carrying out a complex set of actions.
Touchscreens/kiosks for customer interface (Examples: self-checkout, self-check-in, touchscreen ordering)	A computer with a touchscreen that allows a customer to receive information or perform tasks related to the business such as registering for a service or purchasing items.
Voice recognition software	Software that converts speech to text or executes simple commands based on a limited vocabulary or executes more complex commands when combined with natural language processing.

Notes: "AI-related" technologies for the purposes of this study are in bold. Definitions validated during cognitive testing of the survey instrument. See Zolas et al. (2020) for further details.

Table 3: Top Non-manufacturing Industries (4-digit NAICS) for AI Use

NAICS	NAICS Description	Mean (All Industries)	0.058
6215	Medical and Diagnostic Laboratories		0.237
6211	Offices of Physicians		0.190
5112	Software Publishers		0.156
5415	Computer Systems Design and Related Services		0.141
5182	Data Processing, Hosting, and Related Services		0.139
5191	Other Information Services		0.126
6214	Outpatient Care Centers		0.121
5411	Legal Services		0.120
1151	Support Activities for Crop Production		0.119
4248	Beer, Wine, and Distilled Alcoholic Beverage Merchant Wholesalers		0.112

Notes: Tabulated from the ABS–LBD linked sample (column 1, Table 1). Shares computed using the LBD tabulation weights of firm counts, divided by the total number of firms (including those that responded with "Don't Know" or missing), then scaled up by the total number of non-missing and "Don't Know" responses for each technology. Industry tabulations for multi-unit firms are generated from the largest payroll industry within the firm (if there is a tie, then the industry with the most employment is used).

Table 4: Correlates of AI Use - Firm Size, Age, and Other Technology Use

Description	**(1) Use AI**	**(2) Use AI**
Emp. pctile 51–75	0.011***	0.004***
	(0.001)	(0.001)
Emp. pctile 76–90	0.025***	0.012***
	(0.001)	(0.001)
Emp. pctile 91–95	0.036***	0.019***
	(0.002)	(0.002)
Emp. pctile 96–99	0.053***	0.032***
	(0.002)	(0.002)
Emp. pctile 99+	0.086***	0.063***
	(0.004)	(0.004)
Age pctile 51–75	-0.006***	-0.001
	(0.001)	(0.001)
Age pctile 76–90	-0.011***	-0.004***
	(0.001)	(0.001)
Age pctile 91–95	-0.015***	-0.006***
	(0.002)	(0.002)
Age pctile 96–99	-0.022***	-0.013***
	(0.003)	(0.003)
Age pctile 99+	-0.008***	-0.001
	(0.002)	(0.002)
Use: Digitization		0.019***
		(0.001)
Use: Cloud		0.050***
		(0.001)
State-by-Ind Controls	Yes	Yes
Observations (rounded)	447,000	447,000
R-squared	0.102	0.115

Notes: Standard errors in parentheses. *, ** and *** denotes significance at the 10, 5 and 1% respectively. The coefficient estimates represent the relative AI usage rates across size and age percentiles within a 6-digit (NAICS) industry, by state (State-by-Ind controls). The underlying sample is the baseline sample (column 2, Table 1).

Table 5: Firm Characteristics by AI Use Status (%)
Startup Sample

	(1) Full Sample	(2) Use AI	(3) Do not use AI	(4) Difference
A. Primary Owner Characteristics				
Hold Advanced Degree	22.4	31.5	21.9	9.6
Owned Prior Business	41.4	43.8	41.2	2.6
Missing Age	7.4	9.4	7.2	2.2
Owner Age (0–34)	11.5	11.2	11.5	-0.3
Owner Age (35–54)	54.5	55.1	54.5	0.6
Owner Age (55+)	26.6	24.7	26.8	-2.1
"Very Important" Reasons for Owning the Business				
Wanted to be own boss	63.9	62.9	64	-1.1
Flexible hours	54.5	52.8	54.7	-1.9
Balance work and family	56.4	55.1	56.4	-1.3
Opportunity for greater income	63.0	62.9	63	-0.1
Best avenue for ideas/goods/services	56.9	61.8	56.7	5.1
Unable to find employment	7.2	9.0	7.1	1.9
Working for someone else not appealing	30.5	31.5	30.5	1.0
Always wanted to start own business	48.5	49.4	48.5	0.9
Entrepreneurial role model	24.5	27.0	24.3	2.7
Carry on family business	9.9	11.5	9.9	1.6
Help or become involved in community	24.1	30.3	23.7	6.6
Other reasons	8.0	10.6	7.8	2.8
Lifestyle reason	63.6	62.9	63.7	-0.8
B. Startup Financing				
Funded by VC	0.9	2.9	0.8	2.1
Missing Startup Capitalization	8.4	6.3	8.5	-2.2
Startup Capitalization <25K	38.4	38.2	38.4	-0.2
Startup Capitalization 25K–1M	38.9	42.7	38.8	3.9
Startup Capitalization 1M+	2.5	3.8	2.4	1.4
Don't Know Startup Capitalization	11.8	10.3	11.9	-1.6
C. Innovation & Business Strategy				
Process Innovation	19.9	39.3	18.7	20.6
Product Innovation	53.6	66.3	52.8	13.5
Patents Owned or Pending	2.1	5.2	1.9	3.3
IP is very important	19.5	40.4	18.1	22.3
Growth-oriented innovation strategy	69.5	77.5	68.9	8.6
D. Geography				
In Micropolitan or Rural CBSA ($< 50,000$)	22.8	20.0	23.0	-3.0
In Small-Sized CBSA ($< 250,000$)	9.1	8.1	9.1	-1.0
In Medium-Sized CBSA ($< 1,000,000$)	17.8	17.1	17.8	-0.7
In Large-Sized CBSA ($1,000,000+$)	50.3	55.1	49.9	5.2

Notes: "Use" is defined as having responded with "In use for less than 5% of production or service", "In use for between 5%–25% of production or service" or "In use for more than 25% of production or service" for the category listed on any of the AI-based Business Technologies. The "Lifestyle reason" indicator is equal to 1 if the primary owner responded that either "Flexible hours" or "Balance work and family" was a very important reason for owning the business. The "Growth-oriented innovation strategy" indicator is equal to 1 if a firm responded that a focus on introducing new goods or services, reaching new customer groups, or opening up new domestic or export markets was a "very important" innovation strategy for the business. The underlying sample is our startup sample (column 4, Table 1).

Table 6a: Correlates of Early Firm Growth

Description	(1) Revenue Growth (First 3 years)	(2) Revenue Growth (First 3 years)	(3) Revenue Growth (First 3 years)	(4) Revenue Growth (First 3 years)	(5) Revenue Growth (First 3 years)	(6) Specific Industries Revenue Growth (First 3 years)
Advanced Degree (1/0)	0.0248***			0.0161**	0.0167**	0.0129
	(0.0069)			(0.0068)	(0.0068)	(0.0108)
Prior Business (1/0)	0.0194***			0.0094**	0.0076*	0.0107
	(0.0045)			(0.0045)	(0.0045)	(0.0079)
Owner Age (35-54)	-0.0463***			-0.0446***	-0.0428***	-0.0445***
	(0.0070)			(0.0070)	(0.0070)	(0.0134)
Owner Age (55+)	-0.1009***			-0.0961***	-0.0937***	-0.1037***
	(0.0079)			(0.0079)	(0.0079)	(0.0151)
Lifestyle Reason (1/0)	-0.0037			-0.0081*		
	(0.0049)			(0.0049)		
Funded by Venture Capital (1/0)		0.1607***		0.1289***	0.1276***	0.1897***
		(0.0364)		(0.0362)	(0.0362)	(0.0633)
Startup Capitalization 25K-1M (1/0)		0.0629***		0.0564***	0.0553***	0.0965***
		(0.0051)		(0.0051)	(0.0051)	(0.0095)
Startup Capitalization 1M+ (1/0)		0.1393***		0.1308***	0.1295***	0.2109***
		(0.0173)		(0.0172)	(0.0172)	(0.0370)
Process Innovation (1/0)			0.0397***	0.0352***	0.0328***	0.0290***
			(0.0056)	(0.0056)	(0.0056)	(0.0096)
Product Innovation (1/0)			0.0132***	0.0100**	0.0093**	0.0157**
			(0.0043)	(0.0043)	(0.0043)	(0.0078)
Patents Owned or Pending (1/0)			0.0773***	0.0591***	0.0586***	0.0621**
			(0.0207)	(0.0205)	(0.0204)	(0.0312)
IP is very important (1/0)			0.0402***	0.0327***	0.0302***	0.0327***
			(0.0061)	(0.0061)	(0.0061)	(0.0105)
Growth-oriented innovation strategy (1/0)			0.0542***	0.0480***	0.0402***	0.0368***
			(0.0047)	(0.0047)	(0.0048)	(0.0084)
State-by-Ind Controls	Yes	Yes	Yes	Yes	Yes	Yes
Firm Age Controls	Yes	Yes	Yes	Yes	Yes	Yes
Owner Gender Controls	Yes	Yes	Yes	Yes	Yes	Yes
Reasons for Owning Controls	No	No	No	No	Yes	Yes
Observations (rounded)	75,000	75,000	75,000	75,000	75,000	28,000
R-squared	0.227	0.229	0.230	0.237	0.238	0.228

Notes: Standard errors in parentheses. *, ** and *** denotes significance at the 10, 5 and 1% respectively. Coefficients are estimated using ordinary least squares (OLS). Revenue growth refers to the three-year average of the log-difference measure of annual revenue growth. State-by-Ind controls refers to state and 6-digit NAICS industry dummies. The underlying sample for Columns 1–5 is our startup sample (column 4, Table 1). Specific Industries in Column 6 refers to the startup sample in specific industries (column 5, Table 1). For a detailed breakdown of "Reasons for Owning Controls" (which replaces "Lifestyle Reason" in Columns 5 and 6), please refer to Table 6b.

Table 6b: Correlates of Early Firm Growth
(Detailed Reasons for Owning Business)
Corresponds to Columns 5 and 6 in Table 6a

Description	(1) Revenue Growth (First 3 years)	(2) Spec. Ind. Revenue Growth (First 3 years)
Wanted to be own boss (1/0)	0.0199***	0.0291***
	(0.0062)	(0.0108)
Flexible hours (1/0)	-0.0175***	-0.0303***
	(0.0061)	(0.0108)
Balance work and family (1/0)	-0.0115*	-0.0103
	(0.0060)	(0.0106)
Opportunity for greater income (1/0)	-0.0002	0.0015
	(0.0058)	(0.0094)
Best avenue for ideas/goods/services (1/0)	0.0139***	0.0127
	(0.0052)	(0.0091)
Unable to find employment (1/0)	-0.0252***	-0.0317**
	(0.0081)	(0.0158)
Working for someone else not appealing (1/0)	-0.0006	-0.0052
	(0.0051)	(0.0092)
Always wanted to start own business (1/0)	0.0105*	0.0173*
	(0.0054)	(0.0095)
Entrepreneurial role model (1/0)	-0.0009	-0.0138
	(0.0057)	(0.0107)
Carry on family business (1/0)	-0.0029	0.0086
	(0.0076)	(0.0177)
Help or become involved in community (1/0)	0.0212***	0.0223**
	(0.0055)	(0.0096)
Other reasons (1/0)	0.0141*	0.0108
	(0.0079)	(0.0145)

Notes: This table breaks out the individual "Reasons for Owning Controls" and directly corresponds to Table 6a. Specifically, Column 1 belongs with Column 5 in Table 6a and Column 2 belongs with Column 6 in the same table. Standard errors in parentheses. *, ** and *** denotes significance at the 10, 5 and 1% respectively. Coefficients are estimated using ordinary least squares (OLS). Revenue growth refers to the three-year average of the log-difference measure of annual revenue growth. State-by-Ind controls refers to state and 6-digit NAICS industry dummies. The underlying sample for Columns 1–5 is our startup sample (column 4, Table 1). Specific Industries in Column 6 refers to the startup sample in specific industries (column 5, Table 1).

Table 7: Correlates of AI Adoption – including Early Firm Growth Rate

Description	(1) Use AI	(2) Use AI	(3) Use AI	(4) Use AI	(5) Use AI	(6) Specific Industries Use AI	(7) Full Owner & Rev. Sample Use AI
Revenue Growth (First 3 years)	0.0134***	0.0128***	0.0112***	0.0133***	0.0058**	0.0150***	0.0021
	(0.0024)	(0.0024)	(0.0024)	(0.0024)	(0.0024)	(0.0045)	(0.0016)
Revenue Growth (Last 3 years)							0.0041**
							(0.0018)
Advanced Degree (1/0)		0.0104***			0.0043	0.0159***	0.0087***
		(0.0032)			(0.0031)	(0.0054)	(0.0018)
Prior Business (1/0)		0.0120***			0.0066***	0.0081*	0.0050***
		(0.0023)			(0.0023)	(0.0047)	(0.0013)
Owner Age (35-54)		-0.0029			0.0016	0.0046	0.0004
		(0.0034)			(0.0033)	(0.0069)	(0.0026)
Owner Age (55+)		-0.0092**			-0.0019	0.0008	-0.0014
		(0.0038)			(0.0038)	(0.0078)	(0.0027)
Lifestyle Reason (1/0)		-0.0004			-0.0032	-0.0158***	-0.0021
		(0.0024)			(0.0024)	(0.0052)	(0.0013)
Funded by Venture Capital (1/0)			0.1071***		0.0862***	0.1048***	0.0578***
			(0.0161)		(0.0156)	(0.0254)	(0.0101)
Startup Capitalization 25K-1M (1/0)			0.0148***		0.0078***	0.0121**	0.0090***
			(0.0028)		(0.0027)	(0.0058)	(0.0015)
Startup Capitalization 1M+ (1/0)			0.0297***		0.0139*	0.0129	0.0130***
			(0.0078)		(0.0078)	(0.0159)	(0.0044)
Process Innovation (1/0)					0.0544***	0.0614***	0.0574***
					(0.0034)	(0.0062)	(0.0020)
Product Innovation (1/0)					0.0078***	0.0197***	0.0090***
					(0.0022)	(0.0045)	(0.0012)
Patents Owned or Pending (1/0)					0.0353***	0.0361**	0.0329***
					(0.0115)	(0.0160)	(0.0064)
IP is very important (1/0)					0.0605***	0.0656***	0.0572***
					(0.0037)	(0.0063)	(0.0021)
Growth-oriented innovation strategy (1/0)					0.0066***	0.0058	0.0091***
					(0.0023)	(0.0047)	(0.0012)
Small CBSA				-0.0006	-0.0007	-0.0010	0.0003
				(0.0043)	(0.0043)	(0.0101)	(0.0022)
Medium CBSA				0.0003	-0.0010	0.0087	-0.0004
				(0.0036)	(0.0036)	(0.0082)	(0.0019)
Large CBSA				0.0074**	0.0042	0.0075	0.0013
				(0.0031)	(0.0031)	(0.0070)	(0.0017)
State-by-Ind Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm Age Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Owner Gender Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations (rounded)	75,000	75,000	75,000	75,000	75,000	28,000	209,000
R-squared	0.209	0.210	0.211	0.209	0.230	0.228	0.148

Notes: Standard errors in parentheses. *, ** and *** denotes significance at the 10, 5 and 1% respectively. Coefficients are estimated using ordinary least squares (OLS). Revenue growth refers to the three-year average of the log-difference measure of annual revenue growth. "Use" is defined as having responded with "In use for less than 5% of production or service", "In use for between 5%–25% of production or service" or "In use for more than 25% of production or service" for the category listed on any of the AI-based Business Technologies (Automated Guided Vehicles, Natural Language Processing, Machine Learning, Machine Vision and Voice Recognition). State-by-Ind controls refers to state-specific 6-digit NAICS industry dummies. The underlying sample for Columns 1–5 is our startup sample (column 4, Table 1). Specific Industries in Column 6 refers to the startup sample in specific industries (column 5, Table 1). Full Owner & Rev. Sample in Column 7 refers to our owner and revenue sample (column 3, Table 1). Column 7 is identical to column 5 of Table A4, which contains columns similar to this table for the full owner and revenue sample.

Figures

Figure 1: Prevalence of AI-Related Business Technologies (AI) in the United States by 2017

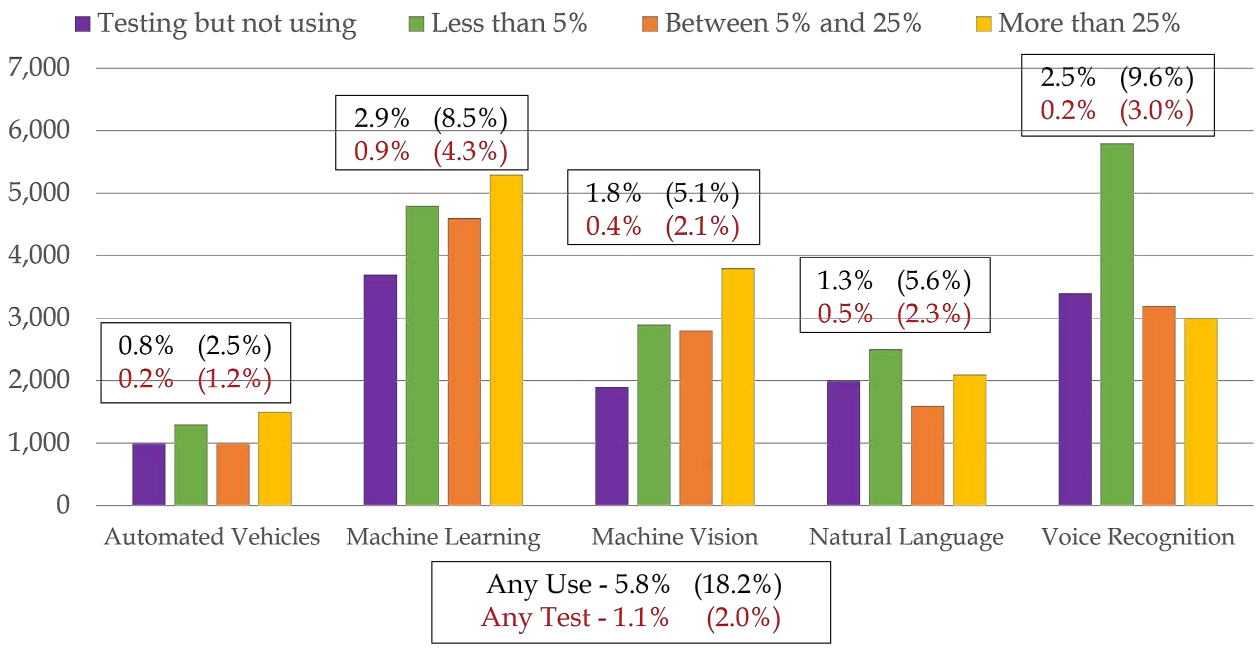

Note: Tabulations (bars) based on unweighted and non-imputed responses from the full ABS–LBD linked sample. "Don't Know" and missing responses are excluded from this figure. The figures listed above the bars reflect the firm-weighted usage rates, followed by employment-weighted usage rates in parentheses (both non-imputed and from our baseline sample). The figures in red font indicate the "Testing" rate for each technology. "Use" is defined as having responded with "In use for less than 5% of production or service", "In use for between 5%–25% of production or service" or "In use for more than 25% of production or service" in the "Business Technologies" module of the ABS. Shares are computed using the LBD tabulation weights of firm counts divided by the total number of firms (excluding those that responded with "Don't Know" or missing). Employment weights are calculated by multiplying the LBD tabulation weights by firm employment.

AI Use Intensity and Testing Rates

Figure 2a: By Firm Size

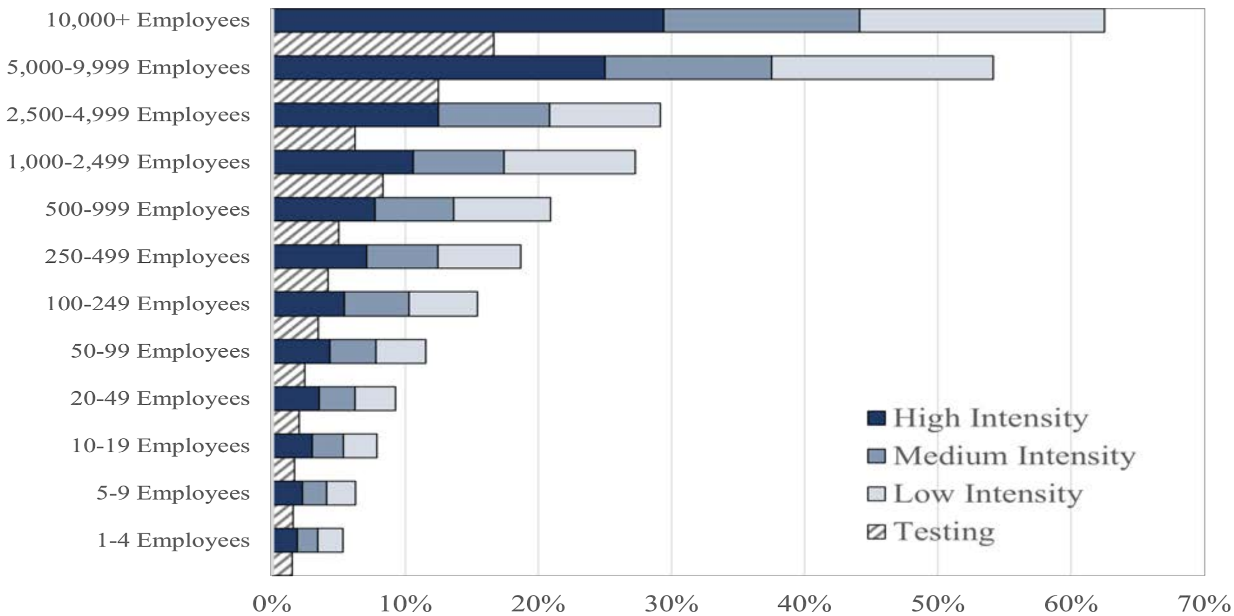

Figure 2b: By Sector

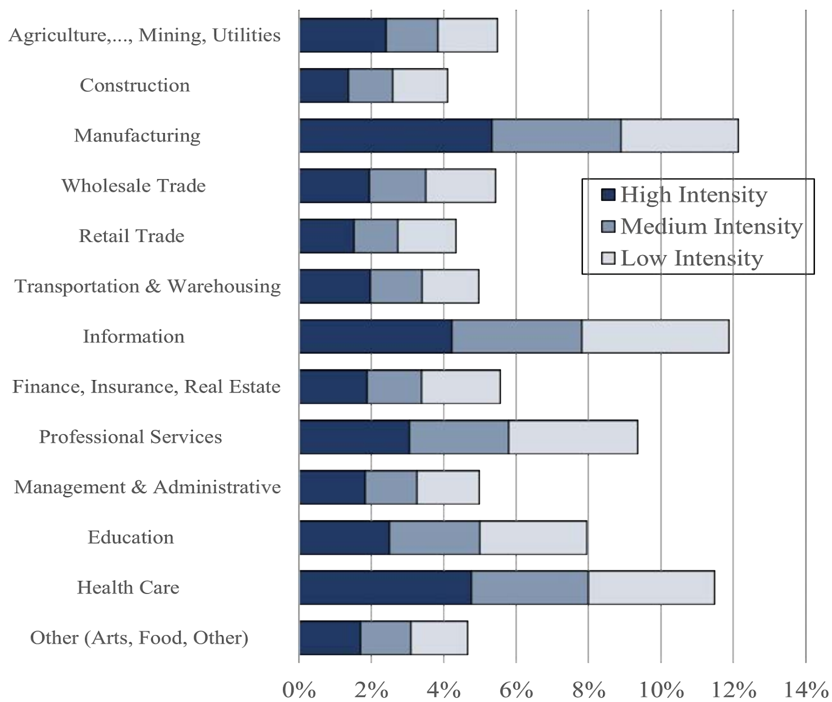

Notes: These figures visually represent the weighted share of firms that indicate intensity of use of at least one of the following business technologies: Automated Guided Vehicles, Machine Learning, Machine Vision, Natural Language Processing, or Voice Recognition (see Table 2) Values are based on the imputed probabilities for respondents who answered "Missing" or "Don't Know" to one or more of the aforementioned business technologies. High intensity corresponds to respondents utilizing at least one of the AI-based business technologies "In use for more than 25% of production or service." Medium intensity corresponds to "In use for between 5%–25% of production or service." Low intensity corresponds to "In use for less than 5% of production or service."

Figure 3: Technology Interdependencies

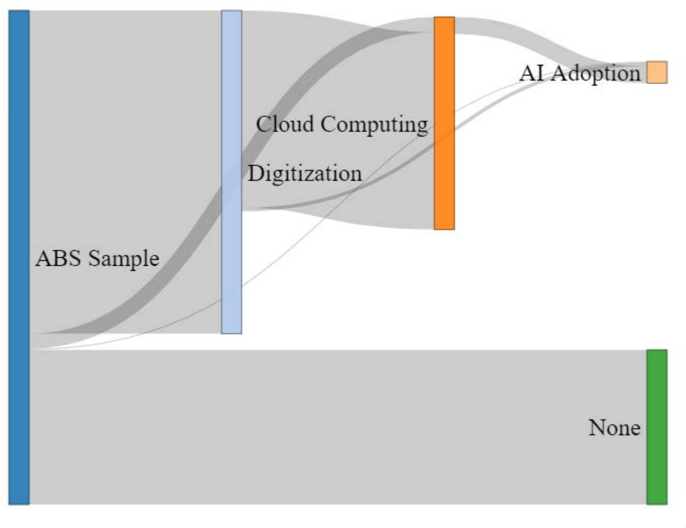

Notes: Sankey representation of firm counts in the ABS–LBD linked sample (column 1, Table 1) as they progress from no technology adoption to reliance on digital information, then cloud computing, and finally AI adoption. The size of the grey area is representative of the number of firm progressing to the next "stage" of technology use. Note that the calculations are made using imputed values for missing responses. Use of AI constitutes the weighted count of firms that indicate use of at least one of the following business technologies: Automated Guided Vehicles, Machine Learning, Machine Vision, Natural Language Processing and Voice Recognition (Table 2).

Figure 4: AI and Robotics Use by Sector

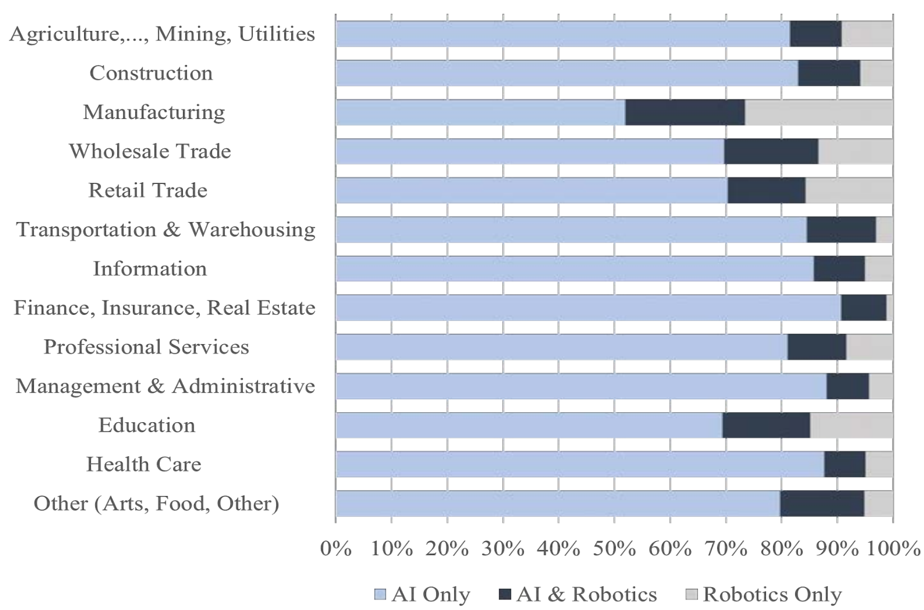

Notes: Percentage of firms that use AI or robotics among firms in the ABS–LBD linked sample that also reported using either an AI-based technology or robotics (roughly 34,000 firms). AI use is defined by the weighted share of firms by sector that indicate use of at least one of the following business technologies: Automated Guided Vehicles, Machine Learning, Machine Vision, Natural Language Processing and Voice Recognition. Robotics use is collected from the same survey question (see Table 2) and similarly weighted.

Figure 5a: AI Use Rates by Single-Unit Startups across Large Core-Based Statistical Areas (CBSAs)

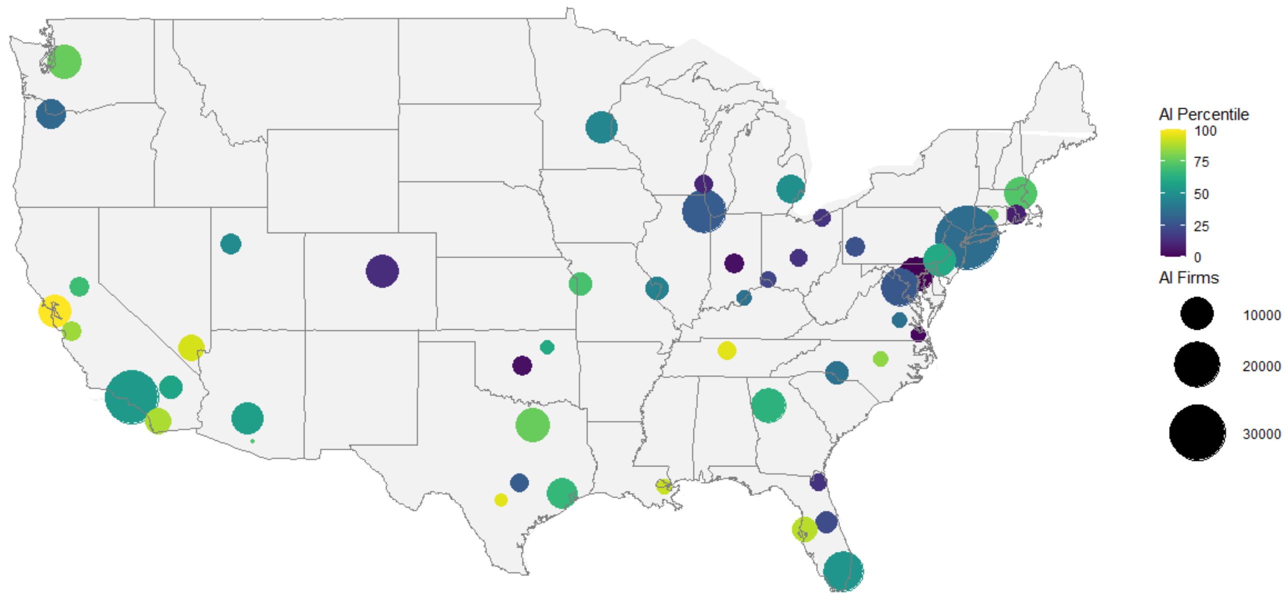

Notes: The map includes the largest (population $> 1,000,000$) Core-Based Statistical Areas (CBSAs). Bubble sizes in the map represent the (firm-weighted) number of single-unit startups (i.e., firm age $\leq$ 5 years) that use AI within each CBSA, and the color gradient represents the percentile rank of the usage rate of AI across single-unit startups within the largest CBSAs. Unlike our "startup sample" described in Table 1 (and used for regressions in Tables 6a, 6b, and 7), the underlying sample of this figure is not restricted to firms with owner or revenue information. Rather, the underlying sample is a subset of our baseline sample—specifically, single-unit startups located in large CBSAs (roughly 30,000 firms).

Figure 5b: AI Use Rates by Young Single-Unit Firms across Large Core-Based Statistical Areas (CBSAs) - Employment-weighted

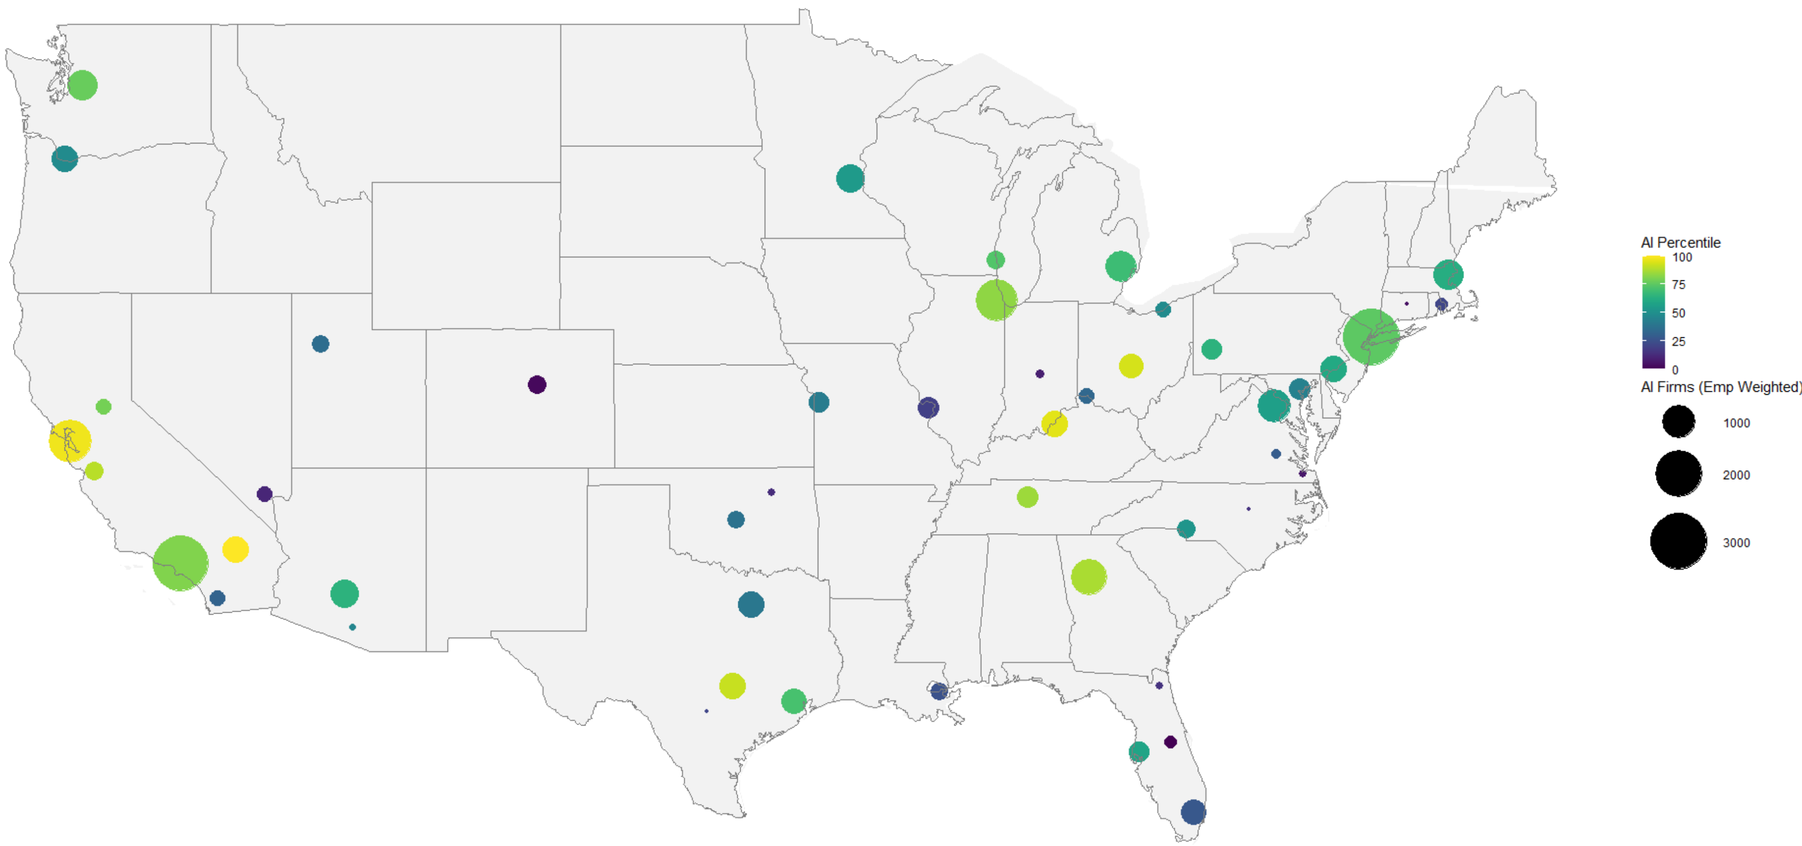

Notes: The map includes the largest (population $> 1,000,000$) Core-Based Statistical Areas (CBSAs). Bubble sizes in the map represent the employment-weighted number of single-unit startups (i.e., firm age ≤ 5 years) that use AI within each CBSA, and the color gradient represents the percentile rank of the usage rate of AI across single-unit startups within the largest CBSAs. Unlike our "startup sample" described in Table 1 (and used for regressions in Tables 6a, 6b, and 7), the underlying sample of this figure is not restricted to firms with owner or revenue information. Rather, the underlying sample is a subset of our baseline sample—specifically, single-unit startups located in large CBSAs (roughly 30,000 firms).

Figure 6a: AI Use Rates by Single-Unit Firms across Core-Based Statistical Areas (CBSAs)

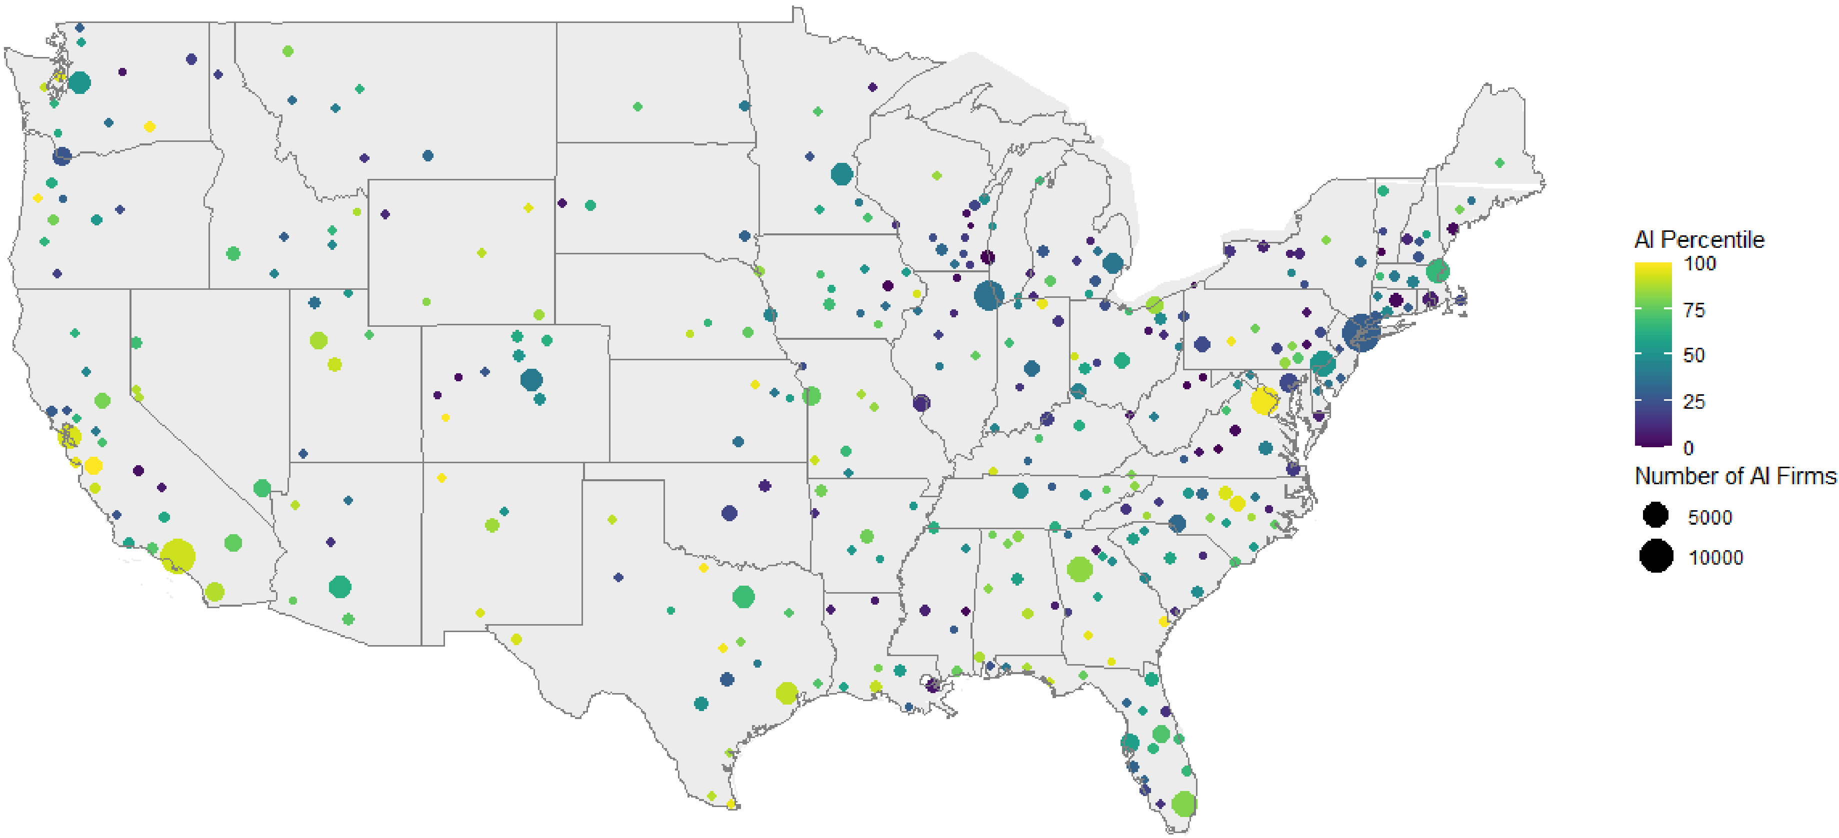

Notes: The map includes the 350 Core-Based Statistical Areas (CBSAs) with at least 20 single-unit firms using AI. Bubble sizes in the map represent the (firm-weighted) number of single-unit firms from our baseline sample that use AI within each CBSA, and the color gradient represents the percentile rank of the adoption rate of AI across single-unit firms within each CBSA type: Micropolitan ($< 50,000$ persons), Small (50,000–250,000 persons), Medium Small (250,000–1,000,000 persons), and Large ($> 1,000,000$ persons).

Figure 6b: AI Use Rates by Single-Unit Firms across Core-Based Statistical Areas (CBSAs) - Employment-weighted

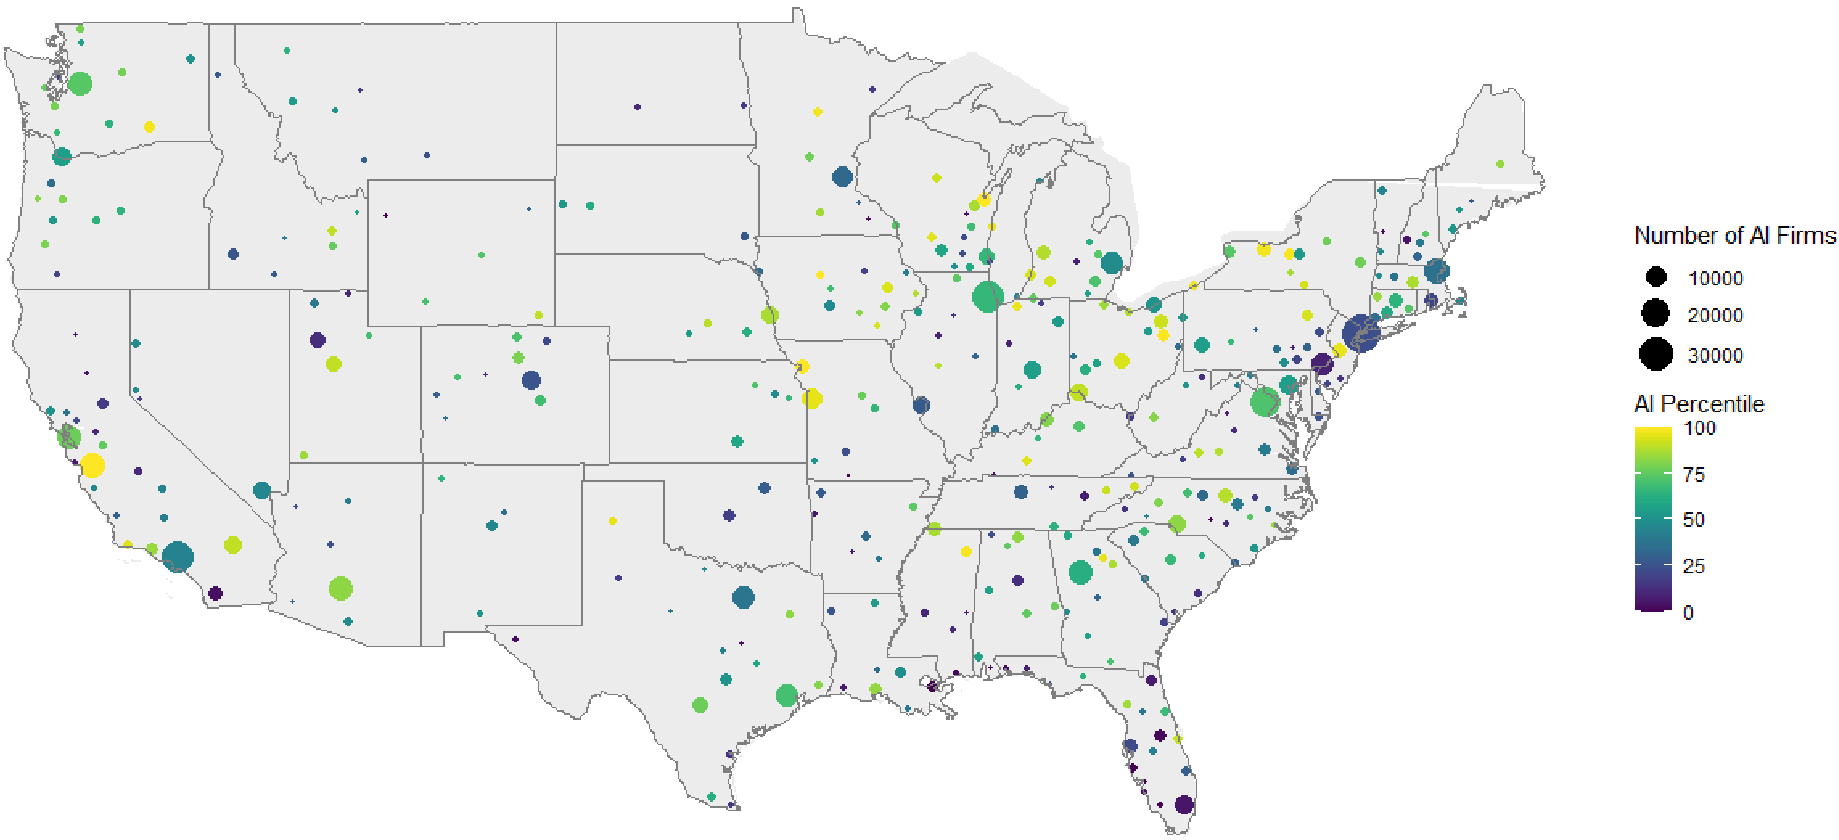

Notes: The map includes the 350 Core-Based Statistical Areas (CBSAs) with at least 20 single-unit firms using AI. Bubble sizes in the map represent the employment-weighted number of single-unit firms from our baseline sample that use AI within each CBSA, and the color gradient represents the percentile rank of the adoption rate of AI across single-unit firms within each CBSA type: Micropolitan ($< 50,000$ persons), Small (50,000–250,000 persons), Medium Small (250,000–1,000,000 persons), and Large ($> 1,000,000$ persons).

Figure 7: Probability of AI Use : Select Marginal Effects (Linear Probability Model)

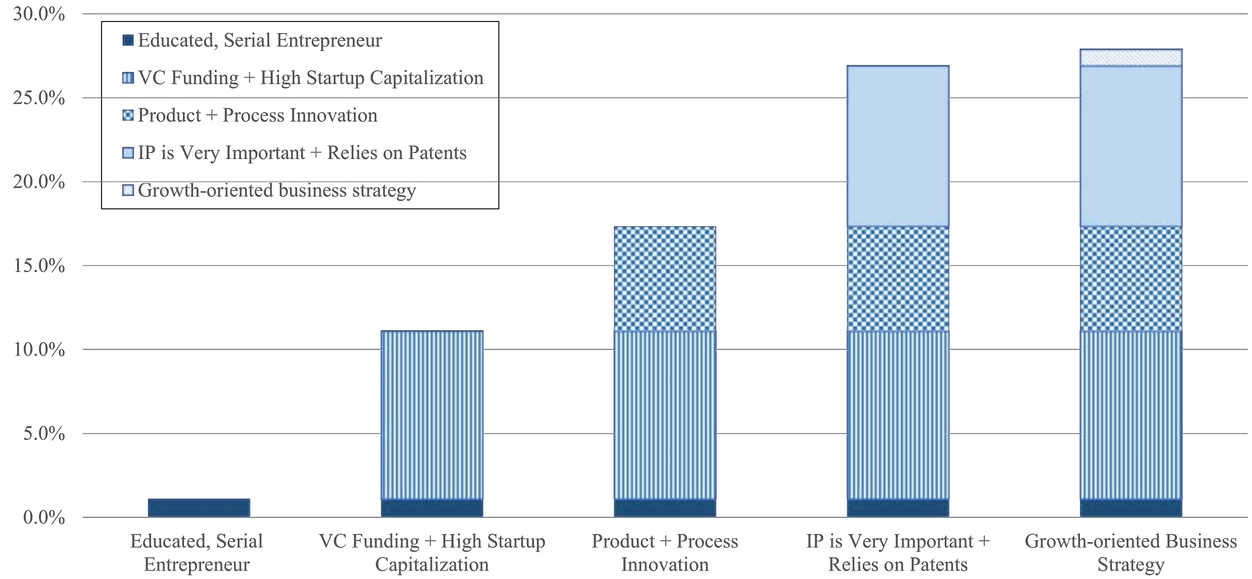

Notes: This figure plots select coefficient values from Column 5 in Table 7 to highlight the cumulative effects of key covariates in determining whether or not a firm adopts AI according to the definition of "use" in Table 7. The variables were grouped as follows: For "Educated, Serial Entrepreneurs", we combined the coefficients for "Advanced Degree" and "Prior Business." For "VC Funding + High Startup Capitalization", we combined the coefficients for "Funded by Venture Capital" and "Startup Capitalization 1M+". For "Product + Process Innovation", we combined the coefficients for "Process Innovation" and "Product Innovation". For "IP is Very Important + Relies on Patents", we combined the coefficients for "IP is very important and "Patents Owned or Pending". Finally, for "Growth-oriented Business Strategy", we combined the coefficient for "Growth-oriented innovation strategy" and took the negative of "Lifestyle Reason". The mean adoption rate is the mean AI usage rate for our startup sample (column 4, Table 1).

A Appendix Tables and Figures

Appendix Tables

Table A1: Firm Characteristics non-manufacturing industries for specific AI Technologies

Automated Guided Vehicles (AGV)			
NAICS	NAICS Description	Mean (All Industries)	0.008
1151	Support Activities for Crop Production		0.074
4245	Farm Product Raw Material Merchant Wholesalers		0.048
2379	Other Heavy and Civil Engineering Construction		0.037
Machine Learning			
NAICS	NAICS Description	Mean (All Industries)	0.029
5112	Software Publishers		0.103
5182	Data Processing, Hosting, and Related Services		0.084
5415	Computer Systems Design and Related Services		0.082
Machine Vision			
NAICS	NAICS Description	Mean (All Industries)	0.018
6115	Technical and Trade Schools		0.048
5112	Software Publishers		0.041
5415	Computer Systems Design and Related Services		0.040
Natural Language Processing			
NAICS	NAICS Description	Mean (All Industries)	0.013
5112	Software Publishers		0.063
5191	Other Information Services		0.059
5415	Computer Systems Design and Related Services		0.056
Voice Recognition			
NAICS	NAICS Description	Mean (All Industries)	0.026
6215	Medical and Diagnostic Laboratories		0.199
6211	Offices of Physicians		0.159
5411	Legal Services		0.093

Notes: Tabulated from the ABS–LBD linked sample (column 1, Table 1).). Shares are computed using the LBD tabulation weights of firm counts, divided by the total number of firms (including those that responded with "Don't Know" or missing). The shares are then scaled up by the total number of non-missing and "Don't Know" responses for each technology. The 2017 industry figures from the LBD are the figures listed in the tables. Industry tabulations for multi-unit firms are generated from the largest payroll industry within the firm (if there is a tie, then the industry with the most employment is used).

Table A2: Large CBSAs by AI Usage Rate among Single-Unit Startups

	(1) Firm-Weighted		(2) Emp.-Weighted	
CBSA	**Share**	**Rank**	**Share**	**Rank**
Atlanta-Sandy Springs-Roswell, GA	6.1	18	12.2	7
Austin-Round Rock, TX	5.4	33	15	5
Baltimore-Columbia-Towson, MD	3	49	6.3	28
Boston-Cambridge-Newton, MA-NH	6.7	13	7.6	19
Charlotte-Concord-Gastonia, NC-SC	5.6	29	7	23
Chicago-Naperville-Elgin, IL-IN-WI	5.4	33	10.3	9
Cincinnati, OH-KY-IN	5	39	5.4	33
Cleveland-Elyria, OH	4.5	40	6.9	25
Columbus, OH	4.5	40	16.1	4
Dallas-Fort Worth-Arlington, TX	6.8	10	6	30
Denver-Aurora-Lakewood, CO	4.5	40	2.4	48
Detroit-Warren-Dearborn, MI	5.7	25	8.1	16
Hartford-West Hartford-East Hartford, CT	6.8	10	2.5	47
Houston-The Woodlands-Sugar Land, TX	6.2	17	8.2	15
Indianapolis-Carmel-Anderson, IN	3.8	45	2.9	45
Jacksonville, FL	4.5	40	3.3	41
Kansas City, MO-KS	6.6	15	6.1	29
Las Vegas-Henderson-Paradise, NV	7.7	4	3.2	43
Los Angeles-Long Beach-Anaheim, CA	5.8	23	9.4	10
Louisville/Jefferson County, KY-IN	5.6	29	18.8	3
Miami-Fort Lauderdale-West Palm Beach, FL	5.8	23	4.2	36
Milwaukee-Waukesha-West Allis, WI	4.2	44	8.3	14
Minneapolis-St. Paul-Bloomington, MN-WI	5.7	25	7	23
Nashville-Davidson–Murfreesboro–Franklin, TN	8.3	2	11.8	8
New Orleans-Metairie, LA	7.4	5	4	37
New York-Newark-Jersey City, NY-NJ-PA	5.5	32	8.6	13
Oklahoma City, OK	3.8	45	5.8	31
Orlando-Kissimmee-Sanford, FL	5.1	37	2.3	49
Philadelphia-Camden-Wilmington, PA-NJ-DE-MD	6	19	7.5	20
Phoenix-Mesa-Scottsdale, AZ	5.9	21	7.8	17
Pittsburgh, PA	5.1	37	7.8	17
Portland-Vancouver-Hillsboro, OR-WA	5.4	33	6.9	25
Providence-Warwick, RI-MA	3.8	45	3.8	38
Raleigh, NC	6.9	9	3.3	41
Richmond, VA	5.6	29	4.3	35
Riverside-San Bernardino-Ontario, CA	5.9	21	20	1
Sacramento–Roseville–Arden-Arcade, CA	6.4	16	9.1	11
St. Louis, MO-IL	5.7	25	3.6	40
Salt Lake City, UT	5.7	25	5.7	32
San Antonio-New Braunfels, TX	8.3	2	3.8	38
San Diego-Carlsbad, CA	7.4	5	5.1	34
San Francisco-Oakland-Hayward, CA	9.5	1	18.9	2
San Jose-Sunnyvale-Santa Clara, CA	7.1	8	13.6	6
Seattle-Tacoma-Bellevue, WA	6.8	10	8.7	12
Tampa-St. Petersburg-Clearwater, FL	7.4	5	7.4	21
Tucson, AZ	6.7	13	6.4	27
Tulsa, OK	6	19	3.2	43
Virginia Beach-Norfolk-Newport News, VA-NC	3.6	48	2.6	46
Washington-Arlington-Alexandria, DC-VA-MD-WV	5.2	36	7.4	21
Mean Micropolitan (< 50,000 persons)	**4.1**	-	**8.2**	-
Mean Small-sized CBSA (50,000 - 250,000 persons)	**4.9**	-	**6.7**	-
Mean Medium-sized CBSA (250,000 - 1,000,000 persons)	**4.5**	-	**6.8**	-
Mean Large-sized CBSA (> 1,000,000 persons)	**5.9**	-	**7.7**	-

Notes: This table reports the AI usage rates among single-unit startups in large CBSAs. Large CBSAs are those with more than 1,000,000 persons. Usage rates are computed using a modified firm weight based on the national share of single-units within strata defined by size class, age class, and 2-digit NAICS sector. Unlike our "startup sample" described in Table 1 (and used for regressions in Tables 6a, 6b, and 7), the underlying sample of this table is not restricted to firms with owner or revenue information. Rather, the underlying sample is a subset of our baseline sample—specifically, single-unit startups located in large CBSAs (roughly 30,000 firms).

Table A3: AI Usage Rates by Firm Characteristics (%)
Startup Sample

	AI Usage Rate (%)		
Full Sample	6.0		-
	Yes	**No**	**Difference**
A. Primary Owner Characteristics			
Hold Advanced Degree	8.4	5.3	3.1
Owned Prior Business	6.4	5.8	0.6
Missing Age	7.7	5.9	1.8
Owner Age (0–34)	5.9	6	-0.1
Owner Age (35–54)	6.1	5.9	0.2
Owner Age (55+)	5.6	6.2	-0.6
"Very Important" Reasons for Owning the Business			
Wanted to be own boss	5.9	6.2	-0.3
Flexible hours	5.8	6.2	-0.4
Balance work and family	5.9	6.2	-0.3
Opportunity for greater income	6	6	0
Best avenue for ideas/goods/services	6.5	5.3	1.2
Unable to find employment	7.5	5.9	1.6
Working for someone else not appealing	6.2	5.9	0.3
Always wanted to start own business	6.1	5.9	0.2
Entrepreneurial role model	6.6	5.8	0.8
Carry on family business	6.9	5.9	1
Help or become involved in community	7.6	5.5	2.1
Other reasons	8	5.8	2.2
Lifestyle reason	5.9	6.1	-0.2
B. Startup Conditions			
Funded by VC	18.8	5.9	12.9
Missing Startup Capitalization	4.5	6.2	-1.7
Startup Capitalization <25K	6	6	0
Startup Capitalization 25K–1M	6.6	5.6	1
Startup Capitalization 1M+	9.2	5.9	3.3
Don't Know Startup Capitalization	5.3	6.1	-0.8
C. Innovative Strategies & Outcomes			
Process Innovation	11.9	4.6	7.3
Product Innovation	7.4	4.4	3
Patents Owned or Pending	14.8	5.8	9
IP is very important	12.5	4.4	8.1
Growth-oriented innovation strategy	6.7	4.4	2.3
D. Geography			
In Micropolitan or Rural CBSA ($< 50,000$)	5.3	6.2	-0.9
In Small-Sized CBSA ($< 250,000$)	5.3	6.1	-0.8
In Medium-Sized CBSA ($< 1,000,000$)	5.8	6.1	-0.3
In Large-Sized CBSA ($1,000,000+$)	6.6	5.4	1.2

Notes: "Use" is defined as having responded with "In use for less than 5% of production or service", "In use for between 5%–25% of production or service" or "In use for more than 25% of production or service" for the category listed on any of the AI-based Business Technologies. The "Lifestyle reason" indicator is equal to 1 if the primary owner responded that either "Flexible hours" or "Balance work and family" was a very important reason for owning the business. The "Growth-oriented innovation strategy" indicator is equal to 1 if a firm responded that a focus on introducing new goods or services, reaching new customer groups, or opening up new domestic or export markets was a "very important" strategy for the business. The underlying sample for this table is our startup sample (column 4, Table 1).

Table A4: Correlates of AI Adoption – including Early and Later Firm Growth Rates Owner & Revenue Sample

Description	**(1) Use AI**	**(2) Use AI**	**(3) Use AI**	**(4) Use AI**	**(5) Use AI**
Revenue Growth (First 3 years)	0.0075***	0.0073***	0.0061***	0.0074***	0.0021
	(0.0016)	(0.0016)	(0.0016)	(0.0016)	(0.0016)
Revenue Growth (Last 3 years)	0.0093***	0.0088***	0.0083***	0.0093***	0.0041**
	(0.0019)	(0.0019)	(0.0019)	(0.0019)	(0.0018)
Advanced Degree (1/0)		0.0148***			0.0087***
		(0.0019)			(0.0018)
Prior Business (1/0)		0.0102***			0.0050***
		(0.0013)			(0.0013)
Owner Age (35-54)		-0.0039			0.0004
		(0.0026)			(0.0026)
Owner Age (55+)		-0.0082***			-0.0014
		(0.0027)			(0.0027)
Lifestyle Reason (1/0)		0.0012			-0.0021
		(0.0013)			(0.0013)
Funded by Venture Capital (1/0)			0.0770***		0.0578***
			(0.0103)		(0.0101)
Startup Capitalization 25K-1M (1/0)			0.0159***		0.0090***
			(0.0015)		(0.0015)
Startup Capitalization 1M+ (1/0)			0.0271***		0.0130***
			(0.0044)		(0.0044)
Process Innovation (1/0)					0.0574***
					(0.0020)
Product Innovation (1/0)					0.0090***
					(0.0012)
Patents Owned or Pending (1/0)					0.0329***
					(0.0064)
IP is very important (1/0)					0.0572***
					(0.0021)
Growth-oriented innovation strategy (1/0)					0.0091***
					(0.0012)
Small CBSA				0.0012	0.0003
				(0.0023)	(0.0022)
Medium CBSA				0.0017	-0.0004
				(0.0019)	(0.0019)
Large CBSA				0.0046***	0.0013
				(0.0017)	(0.0017)
State-by-Ind Controls	Yes	Yes	Yes	Yes	Yes
Firm Age Controls	Yes	Yes	Yes	Yes	Yes
Owner Gender Controls	Yes	Yes	Yes	Yes	Yes
Observations (rounded)	209,000	209,000	209,000	209,000	209,000
R-squared	0.125	0.126	0.127	0.125	0.148

Notes: Standard errors in parentheses. *, ** and *** denotes significance at the 10, 5 and 1% respectively. Revenue growth refers to the three-year average of the log-difference measure of annual revenue growth. "Use" is defined as having responded with "In use for less than 5% of production or service", "In use for between 5%–25% of production or service" or "In use for more than 25% of production or service" for the category listed on any of the AI-based Business Technologies (Automated Guided Vehicles, Natural Language Processing, Machine Learning, Machine Vision and Voice Recognition). State-by-Ind controls refers to state and 6-digit NAICS industry dummies. The underlying sample for this table is our owner and revenue sample (column 3, Table 1). Note that Column 5 is identical to column 7 of Table 7.

Appendix Figures

Figure A1: AI and Robotics Use by Firm Size
Full ABS Sample

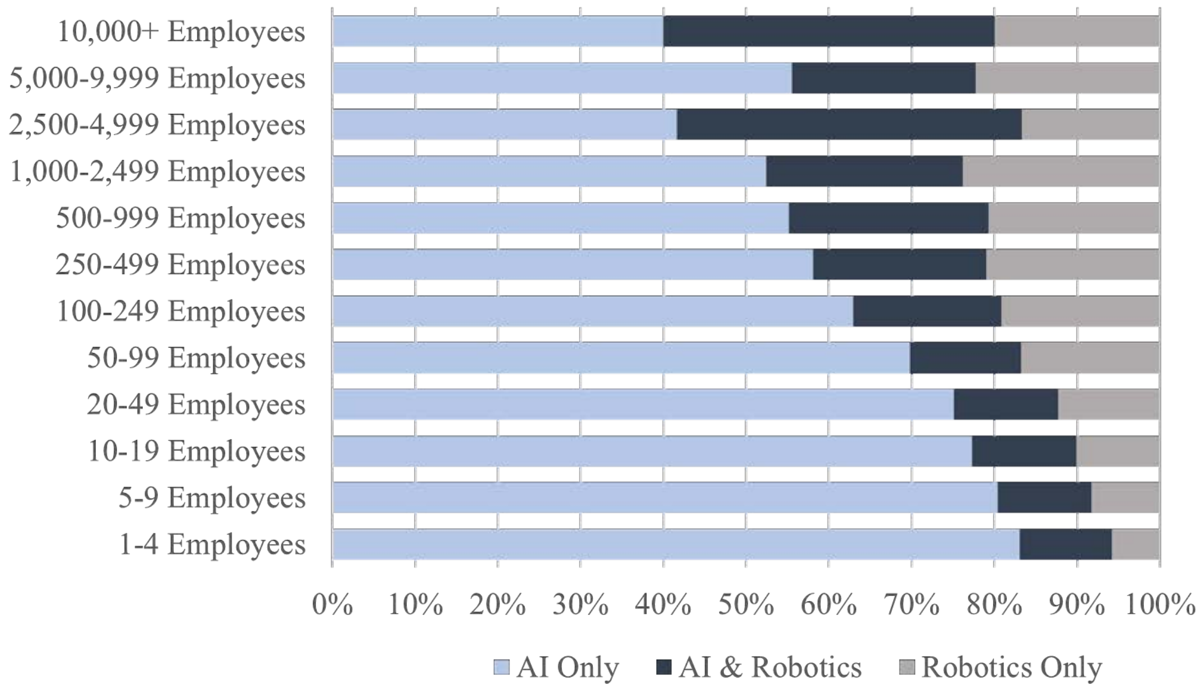

Notes: This figure reports the percentage of firms that use AI or robotics among firms in the ABS–LBD linked sample (column 1, Table 1) that reported using either an AI-based technology or robotics (roughly 34,000 firms). AI use is defined by the weighted share of firms by sector that indicate intensity of use of at least 1 of the following business technologies: Automated Guided Vehicles, Machine Learning, Machine Vision, Natural Language Processing and Voice Recognition. Robotics use is collected from the same survey question (see Table 2) and similarly weighted.

Figure A2: Technological Interdependencies by Intensity

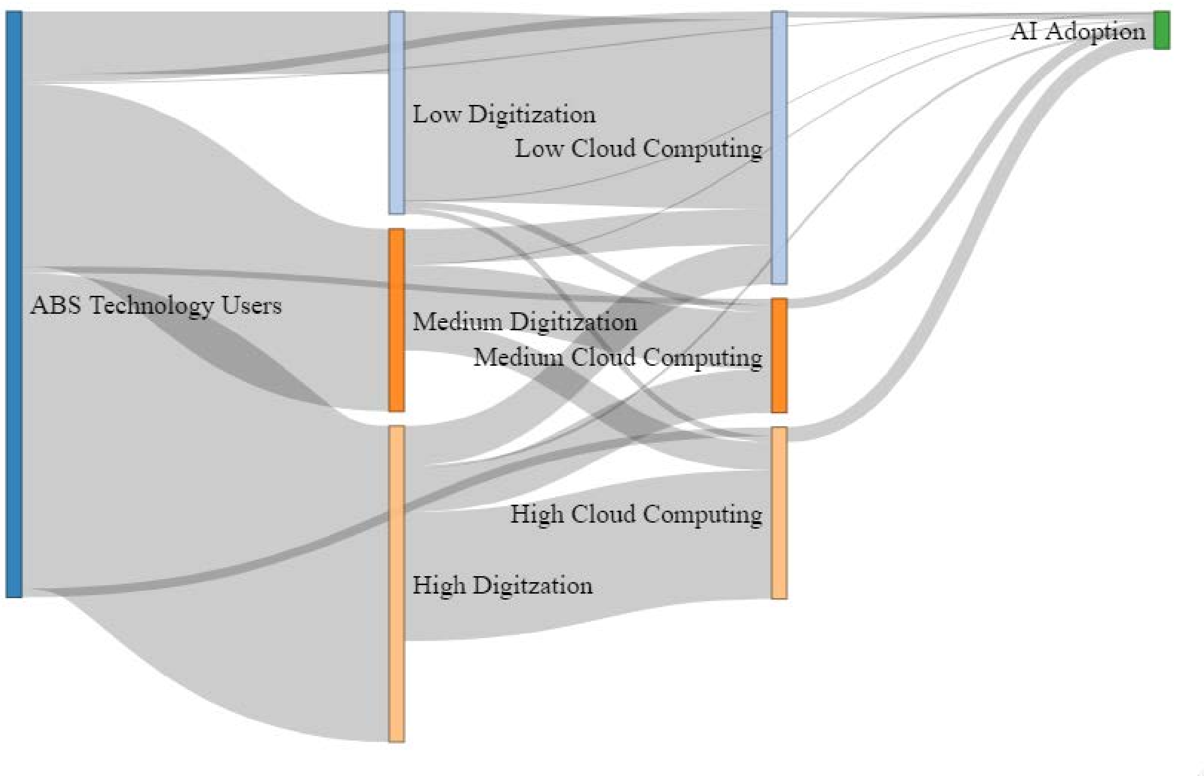

Notes: Sankey representation of firm counts in the ABS–LBD linked sample (column 1, Table 1) as they progress from no technology adoption to reliance on digital information, then cloud computing, and finally AI adoption. The size of the grey area is representative of the number of firm progressing to the next "stage" of technology use. Note that the calculations are made using imputed values for missing responses. Use of AI constitutes the weighted count of firms that indicate use intensity of at least 1 of the following business technologies: Automated Guided Vehicles, Machine Learning, Machine Vision, Natural Language Processing and Voice Recognition (Table 2). Low, Medium, and High intensity use of Digitization and Cloud Computing refer to the respondent answering "All" (corresponds to "High"), "More than 50%" (corresponds to "Medium") and "Less than 50%" corresponds to "Low" for at least one Digital Share of Business Activity or Cloud Service Purchases.

B Data Appendix

B.1 Overview of the ABS

The Annual Business Survey (ABS) was first conducted in 2018 in partnership between the Census Bureau and the National Center for Science and Engineering Statistics (NCSES) within the National Science Foundation. Sent to roughly 850,000 firms nation-wide, it covers all non-farm employer businesses filing the IRS 941, 944, or 1120 tax forms (which includes publicly traded firms). Firms surveyed are selected based on a sampling method that stratifies firms by state, industry, and ownership ethnicity and gender (for details on the 2018 ABS sampling methodology, visit https://www.census.gov/programs-surveys/abs/technical-documentation/methodology.2018.html).

The ABS consolidates three earlier data collections: the Survey of Business Owners (SBO), the Annual Survey of Entrepreneurs (ASE), and the Business R&D and Innovation Survey for Microbusinesses (BRDI-M). The SBO was conducted by the Census Bureau; the ASE was conducted by the Census Bureau in partnership with the Ewing Marion Kauffman Foundation and the Minority Business Development Agency; and the BRDI-M was conducted by the Census Bureau under a partnership with NCSES. The ABS also includes an innovation section—in part derived from the 2016 Business R&D and Innovation Survey (BRDIS)—to provide an expanded set of nationally representative business innovation statistics As a result of these consolidations from earlier collections, the ABS captures many important features of U.S. businesses, including ownership and owner characteristics, startup or acquisition funding, initial capitalization and firm financing, intellectual property strategy, and several aspects of innovation and R&D.

The only purpose-designed questions for this inaugural wave of the ABS pertained to digitization, cloud computing, and advanced business technology use, as described in Section 3. See Zolas et al. (2020) for additional details on this technology module, and see https://www.census.gov/programs-surveys/abs/about.html for more details about the ABS overall.

B.2 Linking to Administrative Data

Approximately 590,000 firms responded to the survey, and about 573,000 were linked to the Longitudinal Business Database (LBD) based on identifiers maintained by the US Census Bureau. This is the leading source of firm-level administrate data on employment and payroll in the United States. Revenue data became available in the LBD as of 1997. Thus, all analyses relying on revenue information are restricted to firms that a) survived through 2018 and b) were 20 years old or younger as of 2017.

B.3 Measuring Owner Characteristics, Startup Financing, and Firm Innovation and Business Strategies in the ABS

While our detailed analysis of the organizational context of AI use in this study focuses on startups (i.e., firms 5 years old or younger), such details are available for a much broader sample of firms. Where useful, we expand our sample from the core "startup sample" of 75,000 firms to include older firms for which this information is available. This larger subsample, however, has a number of limitations worth noting. In particular, it is restricted to firms with information on the primary owner, thus necessarily excluding any publicly-owned firms, as well as estates, trusts, government and tribal entities, associations, membership clubs, cooperatives and foreign entities. For analysis requiring revenue information (notably, not the mapping exercises), this further restricts the sample to roughly 209,000 firms with both ownership and revue data (see Table 1).

B.3.1 Education

The ABS asks "Prior to establishing, purchasing, or acquiring this business, what was the highest degree or level of school" the owner completed? There are 9 checkbox responses ranging from "Less than high school/secondary school graduate" to "Professional Degree, beyond a Bachelor's Degree (for example, MD, DDS, DVM, LLB, JD)." We construct an indicator of having an advanced degree, defined as having a Master's, Doctoral, or Professional degree.

B.3.2 Serial Entrepreneurship

The ABS asks, "Prior to establishing, purchasing, or acquiring this business, how many previous businesses has the owner owned?" Six possible responses range from 0 through 5 or more, though the important variation is captured by an indicator for any prior ownership.

B.3.3 Owner Age

Six checkboxes solicit owner age: under 25, 25–34, 35–44, 45–54, 55–64, and 65 or over. We collapse this into three informative categories: less than 35 years (omitted group), 35–54 years, and 55+ years.

B.3.4 Owner Aspirations

The ABS asks about how important (very important, somewhat important, and not at all important) are twelve potential reasons for owning this business. These reasons include: "Wanted to be my own boss," "Flexible hours," "Balance work and family," "Opportunity for greater income," "Best avenue for my ideas/goods/services," "Unable to find employment," "Working for someone else didn't appeal to me," "Always wanted to start my own business," "An entrepreneurial friend or family member was my role model," "Wanted to carry on the family business," wanted to "Help and/or become more involved in my community," and "Other (Specify)." We construct an indicator which receives a value of 1 if the primary owner reports either "flexible hours" or "balance work and family" as a "very important" reason for owning the business.

B.3.5 Funding Sources

For sources of capital, there are twelve checkbox options provided and respondents can check all that apply. Options include personal/family sources (separately for savings, other assets, home equity loans, and credit cards), business credit cards, business loans (separately for government-guaranteed, bank or financial institution, and government), investment (separately for from family/friends and venture capital), grants, and other sources.

B.3.6 Funding Levels

Regarding the total amount of capital, ten checkboxes range from less than 5,000 dollars to 3 million dollars or more, including a "Don't Know" option: less than \$5,000; \$5,000–\$9,999; \$10,000–\$24,999; \$25,000–\$49,999; \$50,000–\$99,999; \$100,000–\$249,999; \$250,000–\$999,999; \$1,000,000–\$2,999,999; and \$3,000,000 or more. We collapse these ranges into three broader ranges: less than \$25,000; \$25,000–\$999,999; and \$1,000,000 or more.

B.3.7 Innovation

The ABS asks about process and product innovation, separately, across the prior 3 years (2015–2017). The process innovation question asks whether the business introduced new or significantly improved manufacturing methods; logistics, delivery, or distribution methods; and supporting activities for such processes (such as maintenance systems or operations for purchasing, accounting or computing). The product innovation question asks whether the business introduced new or significantly improved goods or services. Response options to both are "yes," "no," and "not applicable." We construct indicators for each.

To measure the firm's reliance on formal IP, we first combine two questions that capture the number of U.S. patent applications pending and the number of U.S. patents owned by the business by the end of 2017, collapsing them into an indicator of any patents owned or pending.

We also leverage questions on the ABS about the level of importance (very important, somewhat important, and not at all important) of the following six types of intellectual property protection: utility patents (patents for inventions), design patents (patents for appearance), trademarks, copyrights, trade secrets, and nondisclosure agreements.

B.3.8 Business Strategy

In the section on Innovation, the ABS elicits the importance (very important, somewhat important, and not at all important) of fourteen business strategies over 2015–2017. These strategies include: improving existing goods/services, introducing new goods/services, reaching new customer groups, customer-specific solutions, lowering price, reducing costs, satisfying key customers, developing niche/specialized markets, opening up new domestic markets, opening up new export markets, improving internal processes, improving delivery of existing goods/services, improving workforce, and understanding/meeting customer needs. We consider a business as having a "growth-oriented" strategy if it reported that introducing new goods or services, reaching new customer groups, or opening up new domestic or export markets were "very important."

B.4 Tabulation Weights for Geographic Results

Because the samples for Figures 5a, 5b, 6a, 6b, and Table A2 contain only single-unit firms, these differ somewhat from our nationally representative baseline sample described in Section 3.1. We thus modify our firm weights. Instead of generating weights based on the universe of firms in the LBD, we generate weights to represent the universe of single-unit firms in the LBD. The stratification we use for calculating the new weights does not explicitly include geographic information. A more complete dive into the geography of AI diffusion would likely benefit from a more nuanced weighting scheme that directly incorporates geographic information in some manner. Regarding the method of generating weights, our non-parametric weighting scheme ensures that the weighted number of surveyed ABS firms is equal to the number of firms in the same respective strata in the LBD. But this relies on a critical mass of surveyed firms in each stratum to avoid assigning unusually large values as weights, and including geography as an additional stratum would result in very small cells. We thus recommend a parametric weighting approach for researchers aiming to incorporate geography directly into the design of firm weights. For multi-unit firms that are part of the "startup sample" and serve as part of Tables A3 and 7, we assign a single CBSA by taking the maximum employment within a firm-zip code.