Engineering Value: The Returns to Technological Talent and Investments in Artificial Intelligence*

Daniel Rock
U. Pennsylvania, Wharton

September 30, 2021

Abstract

Implementing deep neural networks has historically required large investments in computational assets to run the algorithms, talented researchers and developers to implement the code, large-scale standardized data for an appropriate use case, and a problem to solve that simpler methods failed to handle well. This paper studies to extent to which firms also earn returns to their employees' AI skill investments and what might drive this value capture. Employees with technological skills are highly complementary to the intangible knowledge assets that firms accumulate. Companies signal that they own assets complementary to AI by employing workers with AI skills. Using over 180 million position records and over 52 million skill records from LinkedIn, I build a panel of firm-level skills to measure the market value of exposure to newly available deep learning talent from the open source launch of Google's TensorFlow (a deep learning software package). AI skills are strongly correlated with market value, though variation in AI skills from 2014-2017 does not explain contemporaneous revenue productivity within firms. Using a variety of difference-in-differences specifications, I show that the TensorFlow launch is associated with an approximate market value increase of $11 million per 1 percent increase in AI skills for firms with assets complementary to AI. Increases in the price of installed firm-specific AI complements following the TensorFlow AI skill shock is the likely cause of market valuation increases for AI adopters.

*I am grateful to Erik Brynjolfsson, Prasanna Tambe, Chad Syverson, and Andrew Lo (my thesis committee) for extensive comments and guidance. I am grateful as well to Sarah Bana, Sagit Bar-Gill, Seth Benzell, Lee Branstetter, Carter Braxton, Tim Bresnahan, Avi Collis, Chiara Farronato, Daniel Fehder, Morgan Frank, Jillian Grennan, Lorin Hitt, Jonathan Hersh, Dave Holtz, Zanele Munyikwa, Frank Nagle, Joshua Krieger, Danielle Li, Christian Peukert, Guillaume Saint-Jacques, Daniela Scur, Sebastian Steffen, Prasanna Tambe, Lynn Wu, Sam Zyontz, and numerous seminar participants for helpful suggestions. The LinkedIn Economic Graph Research team (particularly Di Mo, Guy Berger, Stephen Lynch, and Jacqui Barrett) provided excellent feedback and access to LinkedIn data as part of the LinkedIn Economic Graph Research Challenge. The MIT Initiative on the Digital Economy and Brookings Institution provided generous financial support. All errors are my own.

1 Introduction

"If you are looking for a career where your services will be in high demand, you should find something where you provide a scarce, complementary service to something that is getting ubiquitous and cheap. So what's getting ubiquitous and cheap? Data. And what is complementary to data? Analysis." - Hal Varian, Freakonomics Blog 02/25/2008[1]

There has been a sizeable increase in the progress of Artificial Intelligence (AI) capabilities in the past decade, largely driven by breakthroughs in deep learning. Deep learning is a subset class of algorithms within machine learning (ML), a specific variety of AI, and training deep learning models from scratch is often a significant challenge. Modern deep learning models often have millions of parameters. Historically, implementing deep neural networks (DNNs) has required large investments in computational assets to run the algorithms, talented researchers and developers to implement the code, reasonably large-scale standardized data for an appropriate use case, and a problem to solve that simpler methods failed to handle well. This paper focuses directly on the talent complements to AI-related capital investment. The primary research question for this paper is therefore the following: To what extent do firms also earn returns to their employees' AI skill investments and what might drive this value capture? This line of inquiry is neither limited to AI skills nor technology, yet AI skills are newly emergent and specific technological skills are easily tracked in the digital exhaust of online talent databases (Tambe and Hitt, 2012; Tambe, 2014). If firms and workers face perfectly competitive markets, then workers earn the marginal product of their skill investments. The worker might therefore fully capture the surplus associated with learning to use new technologies. Complementary assets are therefore a possibly important channel for employers to gain from employee skillsets.

With so many bottlenecks, the business value of AI was until recently more speculation than reality. But the awakening of AI interest following the remarkable research progress in deep learning algorithms has led to new corporate AI initiatives. There are a number of supply-side tailwinds for AI adoption. Computational and data storage limitations are considerably mitigated by the proliferation of cloud computing technologies, leveling the playing field for technological competition (Ewens et al., 2018; Jin and McElheran, 2017). Firms are investing in data infrastructure and data collection at a rapid clip (Brynjolfsson et al., 2011; Brynjolfsson and McElheran, 2016; Farboodi et al., 2019; Wu et al., 2020; Tambe et al., 2020a). These new economics of data – that data assets are nonrival (Jones and Tonetti, 2020), have high fixed costs of acquisition with low marginal costs of replication, are highly complementary to analytical talent (Abis and Veldkamp, 2020; Bessen et al., 2020) and serve as prerequisites for implementing AI and other advanced technologies (Iansiti and Lakhani, 2020; Zolas et al., 2021) – motivate new shifts in competition over digital capabilities (Sambamurthy et al., 2003; Bessen, 2020). One by one, the barriers to generating business value from the deep learning variant of AI are falling. AI-related capital investments are therefore becoming more common.

[1] https://freakonomics.com/2008/02/25/hal-varian-answers-your-questions/

The set of mechanisms by which technology workers might generate market value is generally applicable to all kinds of human capital. However, technological skill investments can change or depreciate much faster than other kinds of skills (Deming and Noray, 2018; Horton and Tambe, 2019) . What makes technology workers, and engineers in particular, useful for understanding the underlying value creation processes of workers in firms is this capacity for discrete changes in the competitive environment. Technological shifts therefore supply outside researchers with a chance to study the outcomes for human capital-intensive firms. Studying technological changes can supply insight into how companies and employees both gain from business activity. Ordinarily it is a substantial challenge to look within the firm with granular information about specific types of employed workers and the skills they have. This study is among the first to normalize and deploy detailed data on firm employment over time and how workers contribute to the value of their employers.

This paper connects technological human capital acquisition decisions made by employees to the market values of their employers. I find evidence that the expected future proliferation of AI talent causes previously sidelined AI projects become profitable and existing AI projects become even more profitable ("price effects"). The principal finding is that these price effects at the outset of a skill proliferation event are the most likely mechanism and not contemporaneous productivity increases or overall firm-level worker exposure to AI. These effects increase the value of installed capital that is complementary to AI talent.

Hal Varian's advice in the quote above that one should "find something where you provide a scarce, complementary service to something that is getting ubiquitous and cheap" is equally relevant to firms as costs to acquire the computational and data inputs to ML applications drop. In the current market environment, there are few technologies with the transformative potential of artificial intelligence and machine learning (Agrawal et al., 2018; Brynjolfsson et al., 2018c; Cockburn et al., 2018). Availability of (affordable) talent is nevertheless a remaining challenge for many would-be employers. The returns to investments in ML applications can be hampered by the hunt for AI talent, with top tier scientists earning more than $1 million annually in some cases.[2] In a series of reports, McKinsey research teams reported that by 2016, there were only about 235,000 data scientists in the U.S. labor force (approximately 0.16 percent of workers) (Nicolaus Henke et al., 2016; Bughin et al., 2017). However, between 2013 and 2020, job postings related to machine learning grew from 0.1 percent to 0.5 percent of all postings; postings related to AI grew from 0.03 percent to 0.3 percent (Zhang et al., 2021). What changed? Are all of these firms asking for talent they have no hope of hiring?

In what came as a surprise to many, large information technology firms made their deep learning implementation platforms open source. In particular, Google's decision to open source TensorFlow in November 2015, its internal platform for designing, training, and deploying deep neural networks, led to a staggering increase in the availability of AI talent. By the end of 2015 TensorFlow was already one of the most popular machine learning software libraries available (Zhang et al., 2021).[3]

[2]https://www.nytimes.com/2018/04/19/technology/artificial-intelligence-salaries-openai.html

[3]As measured by GitHub stars, a means of tracking open source engagement on the GitHub software versioning

Given that it was broadly unexpected for Google to make some of its core technology readily available to the machine learning community, the open source launch of TensorFlow constitutes a natural experiment to understand the talent acquisition and valuation effects for firms with previously sunk investments in AI complements. Furthermore by tracking accumulation of skill-level corporate human capital, we can better understand valuation changes due to technological labor exposure. Prior to TensorFlow, the ability to train neural networks was rare and highly specialized. The launch of this tool both effectively commodified deep learning as a skill and accelerated forward expectations for how soon deep learning would be easy to learn more generally. Employers then indeed have an opportunity to provide a scarce complement to an increasingly ubiquitous AI skillset.

I explore how prior investment in AI-related skills at the corporate level facilitated employer value capture as a result of an improved outlook for AI talent markets (primarily via TensorFlow). I define firms with AI complements as those companies which have signaled the presence of assets that make AI productive via observable hiring of workers with AI skillsets. I construct a dataset of matched employer-employee skills stocks and employment for a panel of publicly traded firms using LinkedIn's resume database. Following the introduction of TensorFlow, I find a rapid increase in the rate and quantity of addition of Artificial Intelligence skills on LinkedIn. Using a series of difference-in-difference designs with the firm as the unit of analysis, I find that the TensorFlow shock had differential effects on market value. The value of companies making investments in AI grew more following TensorFlow, even controlling for a wide variety of other complementary skills and including firm fixed effects. For firms in the third and fourth quintile of AI skills, each additional 1% in Artificial Intelligence (AI) skill record counts on LinkedIn is correlated with an increase in firm market value of nearly $3.3 million following the introduction of TensorFlow. Using a synthetic difference-in-difference model, I find that having deep learning talent at the time of the TensorFlow launch provided exposure to increase firm market capitalization by approximately $14 billion in adopter firms.

I test several mechanisms, finding evidence consistent with an increase in the price of AI complement assets and not through productivity enhancements or rapid AI-related asset quantity accumulation following the TensorFlow launch. I further refine the possible source of the price effect to ascertain whether the revalued intangible asset might be firm-level opportunities to apply machine learning causing the increase in market value using the Suitability-for-Machine Learning (SML) measures in (Brynjolfsson and Mitchell, 2017; Brynjolfsson et al., 2018b). Higher average firm exposure to ML-suitable tasks is negatively correlated with market value in the post period, but the timing of negative valuation changes correlated with SML does not coincide with the TensorFlow launch. Similar tests on a number of related placebo skills do not show similar responses to the TensorFlow launch. I also apply alternative specifications within the set of difference-in-difference approaches (including synthetic difference-in-differences (Arkhangelsky et al., 2019)) to test the robustness of the main result. Taken together, these results suggest an upward adjustment in the

platform, than Theano, Caffe, and even Scikit-Learn (a popular ML library) market expectations of future yields for firms that were early to AI investment when the AI skillset became cheaper to acquire.

The paper is organized as follows: Section 2 describes the relevant literature on AI skills and technological investments as well as the TensorFlow context. Section 3 details a stylized theoretical model of how employee skill acquisition can enter the market valuation of employers. Section 4 describes the LinkedIn data and includes summary statistics. Section 5 reports main results across different specifications. Section 6 tests some candidate mechanisms and discusses threats to identification assumptions. Section 7 contains a series of robustness checks, with discussion in Section 8. Section 9 concludes.

2 Context for a Changing Landscape of AI Skills

2.1 Related Research on AI Skills and Technological Investments

Like information technology in general, artificial intelligence-related assets are often intangible, but the returns for AI are mostly in the future at this point (Brynjolfsson et al., 2018c). The recent progress in AI is predominantly a result of advances in deep learning techniques, a specific kind of machine learning approach. Deep learning and neural net algorithms are decades old, but have only recently grown in popularity as large-scale datasets and cheap computational power have made them viable in new domains (White and Rosenblatt, 1963; Rumelhart et al., 1986; LeCun et al., 1998, 2015). As a new kind of software, however, deep learning and AI more broadly appears to be an early-stage general purpose technology (Bresnahan and Trajtenberg, 1995; Bresnahan, 2010; Goldfarb et al., 2020). AI is potentially pervasive, improves over time as better and more data arrive, and can spawn complementary innovation. Within AI, deep learning is also a prediction technology (Agrawal et al., 2017, 2018). Of course, deep learning is not the only prediction technology of its kind – similar problems might be solved by simpler methods like linear regression. Yet the performance of deep learning on formerly insurmountable tasks (e.g. image and speech recognition) has marked a watershed moment in the cost of prediction.

AI capabilities are a type of information technology investment, and IT investments generally necessitate coinvention that leads to an accumulation of intangible assets (Bresnahan et al., 1996; Melville et al., 2004). These intangibles include knowhow, business processes, corporate culture, and organizational designs that allow the new technology to increase corporate productivity. Understandably as a result, intangible assets are also complementary to and correlated with measures of investment in IT categories as well as technological human capital (Bresnahan et al., 2002; Brynjolfsson et al., 2002; Saunders and Tambe, 2015; Saunders and Brynjolfsson, 2016). Further, the shift toward intangible assets in the digital age has opened up a research agenda into the productivity effects of new varieties of IT capital, with technology diffusion serving as a leading explanation for the widening productivity differences between firms at the frontier and firms at the median productivity level (Lustig et al., 2011; Andrews et al., 2015; Tambe et al., 2020a). Intangible assets are inherently hard to measure. Typical approaches to measurement involve monitoring more easily measured complementary investments. Findings following this approach suggest that intangibles constitute an increasingly large component of the U.S. economy's asset stock, and are one possible driver of industrial concentration (Corrado et al., 2009; Marrano et al., 2009; McGrattan, 2017; Haskel and Westlake, 2017; Crouzet and Eberly, 2018, 2019; Bessen, 2020). Firms often fail to capitalize software expenses, for example, but market value of hidden investments is recoverable using observable complements (like measures of labor inputs).[4][5] IT capabilities are also not only cost reducers, but sources of revenue and productivity growth (Dewan et al., 1998; Dewan and Kraemer, 2000; Mithas et al., 2012). This firm-specific set of capabilities is therefore a possible source of rents for IT-intensive companies (Bharadwaj, 2000; Barua et al., 2004). Often these firm-specific returns arise as firms discover unique complementarities between IT and other functions (Barua et al., 1996; Tambe, 2014).

Since prediction is pervasive throughout the economy, the promise of this new AI is that it will lead to (often firm-specific) business process innovation, job redesign, automation, and new engineering advances across many domains in the economy (Furman and Seamans, 2018; Brynjolfsson et al., 2018b; Felten et al., 2018; Webb, 2019). Critically, deep learning mitigates the obstacle of Polanyi's Paradox where "we know more than we can tell" (Polanyi, 1966; Autor, 2014). For supervised deep learning models, we measure inputs and outputs. The map between them is learned by the algorithm. Even the relatively brittle, bespoke applications of deep learning could feasibly cause large shifts in labor demand and economic value creation processes (Brynjolfsson et al., 2018a). Like other technologies with the potential to change and automate work, machine learning will differentially impact tasks that are technologically and socially feasible (Autor et al., 2003; Acemoglu and Autor, 2011; Autor and Dorn, 2013; Acemoglu and Restrepo, 2018; Frank et al., 2019). Possibly one driver of corporate valuation then is the capacity to re-engineer tasks using AI technologies. Still, as a nascent and research-intensive field, deep learning and AI-related asset valuations are likely similar to corporate research assets. Research and development expenses, of which a substantial component is researcher salaries, are reliably strongly correlated with market value and drivers of patenting and other innovative activity (Hall, 1993, 2006; Bardhan et al., 2013; Peters and Taylor, 2017). Given the GPT potential of AI, the return on successful investment in AI complements is likely to be high as well. This high return includes investments in lowering the costs of implementing AI, as the TensorFlow open source decision did by making deep learning easier to learn (for humans).

In AI, the pecuniary benefits of open source innovation, by revealed preference, outstripped the benefits of private IP for Google in TensorFlow's case. Since one of Google's stated aims in open-sourcing TensorFlow was to increase usability and accessibility of deep learning for engineers throughout the economy,[6] A primary pecuniary benefit of making AI models easier to build is an expected future drop in marginal wage rates for AI-intensive human capital. This paper therefore makes a contribution to the open source software (OSS) literature, linking market value and an open source event. OSS has been linked already to productivity (Nagle, 2019) and overall economic

[4] This quantity is the focal object of study in Tambe et al. (2020a)

[5] See Cummins (2005) for a discussion of assumptions and an alternative method using analyst forecasts.

[6] https://www.wired.com/2015/11/google-open-sources-its-artificial-intelligence-engine/

value (Greenstein and Nagle, 2014; Nagaraj, 2021). This paper therefore describes some of the ramifications of the "birth" of a new skill, though similar exogenous events might also cause skill "death" (Horton and Tambe, 2019).[7]

Shifts in the availability of technological talent therefore cause valuation changes via many channels, including but not limited to: price effects on existing assets, appropriability of human capital, marginal labor productivity, and future innovation opportunities. Some of these changes might impact valuation through relatively fast productivity or factor quantity increases, while others are priced in the future. Since skills are held by the worker, and not explicitly a firm asset, a key question is how easy it is for workers to apply their human capital across different employers. Some forms of human capital are subject to competitive bidding pressure from multiple firms, while the returns to firm-specific investments are more subject to bargaining arrangements in contracting. Labor market frictions might therefore create the right incentives for employers to invest in their workers' human capital (Acemoglu and Pischke, 1998, 1999), especially when training is often on-the-job (Becker, 1962). Technological talent is deployed not only in implementing the production function, but also in building the knowledge and business process intangible technology assets of the firm which facilitate growth. As suggested in (Eisfeldt and Papanikolaou, 2013), the organizational capital contribution to equity returns may directly arise as a consequence of employers having a risk exposure to their top employees starting a competitor with better technology.

There are fixed costs of capital investment that apply to human capital as well, wherein quasi-rents can accrue to firms which have already sunk the necessary recruitment and training expenses required to make an employee productive (Hall, 2001). The process for valuation of labor-related assets (not just AI skills) is functionally identical to the valuation of capital assets. Because the marginal adjustment costs of competitors set the price at which the asset is available, firms will hire capital until the marginal adjustment costs of competitors is equal to the marginal value created with that capital (Tobin's q) (Kaldor, 1966; Tobin and Brainard, 1976; Hayashi, 1982).[8] The difference between the firm's adjustment costs and those of its competitors pin down the excess profit of the firm in the short-run. Other frictions like non-negligible search costs for employer-employee matches, and market power in the labor market can also be sources of temporary profits, yet adjustment costs for labor have been estimated at low values in the past (Hall, 2004, 2017). A number of IT business value studies have taken advantage of Tobin's q or similar constructs as a proxy for documenting the investment returns to technology, relying on the intuitive argument that successful IT systems are costly for competitors to replicate (Bharadwaj et al., 1999; Brynjolfsson and Hitt, 2000; Brynjolfsson et al., 2002; Tambe and Hitt, 2012). When productive, IT-related intangible assets are valued highly because they are so difficult to replicate and to market. It would be difficult to for a company to sell off part of its culture, for example, and managing IT (even for outsourcing) requires investment in

[7]Similar studies of technological tool-based and technological knowledge-based exogenous events have addressed how such changes impact various performance measures for firms and other entities. See, for example Ewens et al. (2018); Jin and McElheran (2017) (Amazon Web Services and cloud), Agrawal et al. (2016) (mathematics), Teodoridis (2017) (Microsoft Kinect), Thompson (2017) (multicore processing), and Zyontz (2018) (CRISPR).

[8]These papers lead to the main market value equation I apply in the empirical section.

firm-specific information capabilities (Mani et al., 2010; Mithas et al., 2011; Fitoussi and Gurbaxani, 2012; Mani et al., 2012).

It can take decades for firms to reconfigure production processes around a new GPT, and while those investments are in progress it may appear that productivity is lagging (Brynjolfsson et al., 2021). With effects that are mostly in the future, market value is one of a handful of measures which is sufficiently forward-looking to account for returns to investment activity in the present day. This study builds and measures firm-level proxies for AI investment up to the start of 2018. Given the interest in AI-related skillsets, a number of papers have already started to track the impact of hiring AI talent, answering the call to use firm-level data in AI studies in Raj and Seamans (2018). Recent work tracking AI investments has also tended to focus on labor stocks and flows, if only because labor-based datasets make it easier to track where firms are investing in AI or ML projects. Large scale online labor datasets like CareerBuilder and LinkedIn have been used in a number of studies linking the business value of IT and productivity gains from technology to measurements of labor flows (Tambe and Hitt, 2012; Tambe, 2014; Benzell et al., 2018a). Alekseeva et al. (2020) and Acemoglu et al. (2020) document a rise in AI job postings, with the latter paper showing relatively little effect on employment or wages overall so far. Babina et al. (2020) use a combination of job postings data and employment profiles to show that AI investing firms have faster growing sales and employment, and that AI investment is more prevalent in large firms (matching many of the results in this paper). Bessen et al. (2020) study AI startups' reliance on data inputs. Additionally, many papers have already found evidence of productivity or other positive effects of AI and other algorithms in businesses across many domains, including finance (Grennan and Michaely, 2019; Fuster et al., 2020), hiring (Cowgill, 2018; Li et al., 2020), news (Claussen et al., 2019), board selection (Erel et al., 2018), and monitoring war destruction from space (Mueller et al., 2020). [9]

Technological labor is a well-established driver of corporate market value, innovation, and productivity (Hall, 1993, 2006; Tambe and Hitt, 2012; Tambe, 2014; Tambe et al., 2020a). Technical knowledge is scarce as well. Engineers, research scientists, information technology workers, and other types of technically-skilled labor must invest for years in school and training to build their technical human capital. Their reward for devoting their creative energies to these pursuits is, in part, higher average wages. Still, not all of the capital gains from applying highly specialized knowledge accrue to workers. There are many ways in which the returns to worker investments in human capital might lead to employer value gains.

This paper is about a firm's choice to invest in human capital for workers both present and future. Google's decision to open source TensorFlow, and later Facebook's analogous choice with

[9] Of course firm-specific investment on the employee's part might give their employer a bargaining advantage, among other factors. Many papers consider monopsony power and frictions leading to it (like covenants to not compete) in greater detail (Marx, 2011; Starr, 2015; Jeffers, 2017; Starr et al., 2017; Azar et al., 2018; Caldwell and Danieli, 2018; Schubert et al., 2020), including for technology workers and researchers in particular (Stern, 2004; Roach and Sauermann, 2010; Kokkodis and Ipeirotis, 2016; Balasubramanian et al., 2018; Kokkodis, 2019; Miric and Ozalp, 2020; Tambe et al., 2020b). This paper does not explicitly consider monopsony power, but rather assumes a firm-level bargaining advantage against AI workers is itself manifested as a kind of intangible asset specific to the firm. The results in this paper therefore nest monopsony as a possibility without relying on market power as a certainty.

PyTorch, were in part investment responses to the difficulty in finding AI workers. By propagating one requisite complement to their own systems, these large firms made their own firm-specific assets more valuable. When other firms had already invested in AI given an expectation that the AI skills would be scarce and/or training costs would be high, offering a simple path forward with more abundant talent effectively lowers the cost of realizing returns on AI-related intangible assets. Engineers, as implementers of technology, are highly complementary to the intangible knowledge assets that firms accumulate. Especially if these intangible asset investments are made under the assumption of higher ongoing engineering and asset servicing costs, a change like an open source event can lead to revaluation corporate projects at more favorable prices. Projects that were previously discarded now generate positive net present value (NPV), and existing high NPV projects become more valuable. [10]

2.2 TensorFlow: An Easier Path to Learn Machine Learning

The open-source launch of Google Brain's TensorFlow machine learning toolkit on November 9, 2015 was a departure from expectations that Google would try to safeguard all of its AI-related intellectual property.[11] The project grew out of a 2011 Google Brain initiative called DistBelief to build and train deep neural nets for research and commercial applications (Abadi et al., 2016).[12] TensorFlow was unique among deep learning modules at the time in that it was designed to serve as a single system that could run on a variety of platforms, ranging from mobile devices to large-scale computational systems with multiple GPUs. Its release meant the wide availability of production-level software packages with greater stability and simplicity than other popular packages at the time (e.g. Theano, Caffe, and Torch). At launch, TensorFlow could be installed as a Python module or in C++, taking advantage of popular programming languages to make deep learning available to as many people as possible. There is now support for R as well.

The package also includes a set of software pipelining tools such as TensorBoard, which helps machine learning engineers visualize the computational graph they have built, and performance tracing which helps track threads as they are processed. At the time, few of the comparable systems (Caffe, Chainer, Theano, and Torch) simultaneously supported symbolic differentiation, was written C++ to facilitate high performance production code, and could easily be mapped to many machines at once. Further, the Python interface and training documentation provided a baseline on which the open-source community could improve. What had been an experts' game was, at least in the near future, going to be something any reasonably talented coder could implement. Soon after, additional

[10] One might think of this particularly mechanism as being similar to the technological version of a home mortgage refinancing process. In this way, when an employer makes a capital investment that is a complement to the employee base's technology skills they are implicitly purchasing a call option on future skill availability.

[11] As noted in Wired (https://www.wired.com/2015/11/google-open-sources-its-artificial-intelligence-engine/): "With TensorFlow, however, the company has changed tack, freely sharing some of its newest—and, indeed, most important—software. Yes, Google open sources parts of its Android mobile operating system and so many other smaller software projects. But this is different. In releasing TensorFlow, Google is open sourcing software that sits at the heart of its empire. 'It's a pretty big shift,' says Dean, who helped build so much of the company's groundbreaking data center software..."

[12] Usually called "deep" when a standard neural net architecture has four or more layers.

abstraction layers like Keras (Chollet, 2015) and PyTorch (Paszke et al., 2017), a Pythonized version of the popular Torch software developed by employees at Facebook, would enter as competitors for TensorFlow. Keras is now integrated with TensorFlow, and other packages like fast.ai have grown popular with users. PyTorch in particular is especially popular within the deep learning research community.

But was the TensorFlow open source decision about talent? Oren Etzioni, a machine learning expert and executive director of the Allen Institute for Artificial Intelligence, at the time stated that Google was trying to "attract developers and new hires to its technology".[13] With new technologies, especially open-source software packages, adoption dynamics and value creation can be highly sensitive to network effects (Hippel and Krogh, 2003; Lakani and Hippel, 2002). In the past decade, many large technology firms have shifted to OSS. Google, Facebook, Microsoft, LinkedIn, Uber, and others have made best-in-class software packages available to the programming community, choosing the open strategy.

One interpretation then is that the TensorFlow open source strategy meant Google could capture more of the rents in the economic applications of machine learning. Another is that their software platform would improve with the benefit of a community of contributors. Indeed, as of early 2021 the TensorFlow GitHub repository has over 150 thousand stars and 80 thousand forks, making it among the most successful machine learning software toolkits in existence.[14] The growth in TensorFlow interest was immediate and explosive at its launch, with interest as measured by GitHub stars eclipsing other popular ML libraries like Scikit-Learn nearly immediately (Zhang et al., 2021). Figure 6, from the 2021 AI Index Report, demonstrates the rapid accelerating of TensorFlow interest (Zhang et al., 2021). Figure 7b shows the Google Trends data for "Deep Learning" searches for the same period. Earlier technologies like MapReduce had eventually been developed in OSS formats that would require Google employees to be retrained. Important as well was Google's growing cloud business, which now includes many machine learning tools. With so many complements to benefit from an active machine learning developer community, the open source choice for Google (and later Facebook) might have had a number of motivations. In any case, the toolkit's success is without doubt. Many programmers can pick up deep learning at lower cost. We now turn to why this change might have enhanced corporate valuations other than Google's.

3 A Model of Firm-Specific Assets with Cheap(er) Learning

3.1 Worker Investment in Skills

We start with a simple case of workers with heterogeneous ability who must decide if they want to invest in a costly new skill that can enhance their wages, a model similar in spirit to Spence (1978) and Bedard (2001), but we will assume that the skill confers value to the employer beyond its signal

[13] https://bits.blogs.nytimes.com/2015/11/09/google-offers-free-software-in-bid-to-gain-an-edge-in-machine-learning/?mtrref=undefined

[14] https://github.com/tensorflow/tensorflow

alone. We further assume that each worker of ability θ earns a wage $w(c_t)$ in each period, where $c_t \in \{0,1\}$ is a binary cost variable set to one when the worker invests in (deep learning) skills and zero otherwise. θ is distributed according to continuous distribution $F(\theta)$ which has support on the interval $\theta \in (0, \infty)$. $F(\theta)$ might, for example, take the form of a lognormal distribution. Then in any period t we define worker utility as:

$$u_t(\theta, c_t) = w_t(c_t) - \frac{c_t}{\theta} \tag{1}$$

Utility at time t is a function of static ability θ, the prevailing wage function $w_t(c_t)$, and the cost decision to invest in skills. Assume that workers with the skill always earn more than workers without: $w_t(1) > w_t(0)\ \forall t$.[15] Then within each period, there is θ_t^* such that the worker is indifferent between investing in the skill or not:

$$\begin{aligned} w_t(1) - \frac{1}{\theta_t^*} &= w_t(0) \\ w_t(1) - w_t(0) &= \frac{1}{\theta_t^*} \end{aligned} \tag{2}$$

In any given period, the difference in wages earned with and without the skill for θ^* ability workers is equal to $\frac{1}{\theta_t^*}$, an index of ability. That means that for $F(\theta^*) = Pr(\theta < \theta^*)$ the share of workers that acquire the skill is $1 - F(\theta_t^*)$ in period t. This gives the primary intuition that wage differentials and ability affect the choice to acquire ML skills. But what if the change in unit costs to invest are not always equal to unity? Extending the problem now to make wage an increasing function of ability θ and to include continuous costs reflecting more intensive skill investment $c_t \in (0, \infty)$ with discount factor for period t equal to δ^t, we have the worker's problem to optimize over their investment choice sequence $\{c_t\}_0^\infty$:

$$\begin{aligned} &\max_{\{c_t\}} \left[\int_{t=0}^{\infty} \delta^t (a_t(c_t))\, dt\right] \\ &\text{subject to} \\ &1) \quad a_{t+1} = \frac{1}{\delta^t} a_t + w_t(c_t, \theta) - \frac{c_t}{\theta}\ \forall t \\ &2) \quad a_0 = 0 \end{aligned} \tag{3}$$

Effectively the worker's problem is to maximize their total savings, which they can then use for consumption.[16] Assets start at zero, are a function of skill investment c_t and are equal to a_t for time t. They must make the choice, based on the prevailing wage function $w_t(c_t, \theta)$ and their abilities θ, whether or not to invest in skills for that period, and skills depreciate fully by the next period.[17]

[15] The continuous version of this skill investment cost criterion is that $\frac{\partial w_t}{\partial c_t} > 0$. We will assume this in the next step.

[16] Modeling consumption explicitly would be more realistic but offers little extra insight for this part of the problem (instead we assume workers save whatever they earn).

[17] We also require the Spence-Mirrlees single crossing condition that $\frac{\partial^2 W}{\partial c_t \partial \theta}(c_t, \theta) > 0$.

This reduces to a relatively simple rule. If the increased wage value is larger than the ability-adjusted marginal cost for worker i, the worker should acquire more of the skill at level c_{it}:

$$\int_{t=0}^{\infty} \delta^t [\int_{c_t=0}^{c_{it}} (\frac{\partial w_t}{\partial c_t})\, dc_t]\, dt \;>\; \int_{t=0}^{\infty} \delta^t \frac{c_{it}}{\theta_i}\, dt \tag{4}$$

and with fixed θ_i over all periods, we get an analog of the above. Worker i should acquire more of the skill if:

$$\frac{\int_{t=0}^{\infty} \delta^t [\int_{c_t=0}^{c_{it}} (\frac{\partial w_t}{\partial c_t})\, dc_t]\, dt}{\int_{t=0}^{\infty} \delta^t c_{it}\, dt} \;>\; \frac{1}{\theta_i} \tag{5}$$

and for each i, it will be the case that $c_{it} = c_{it}^*$ at some level, so the inequality above will hold with equality (the worker is indifferent between learning and not learning the new skill):

$$\frac{\int_{t=0}^{\infty} \delta^t [\int_{c_t=0}^{c_t^*} (\frac{\partial w_t}{\partial c_t})\, dc_t]\, dt}{\int_{t=0}^{\infty} \delta^t c_t^*\, dt} \;=\; \frac{1}{\theta} \tag{6}$$

The numerator is the net present value of the wage differential earned for investing in the skill up to the optimal level, the denominator is the net present value of the unit costs (but not ability adjusted costs) to acquire the skill. For c_{it}^* these will be equivalent, the inequality above will hold with equality. In the binary cost case, the interpretation is the same: increased wages justify investing in deep learning skills, but so do lower costs. It is further possible that for some levels of ability the wage increases are insufficient to justify investment. At this θ^*, anyone with lower ability will not invest at all in the skill.[18] The probability of acquiring the skill overall for a worker with randomly drawn type is $1 - F(\theta^*)$. With N workers in a simple economy, $N(1 - F(\theta^*))$ workers will acquire the skill, albeit at differing levels. Nothing here is specific to machine learning skills, and in reality of course skills probably do not depreciate fully each period. Yet adding a law of motion for skills acquisition preserves the following intuition: we can also see that either 1) raising the present value of the stream of wage premia to learn the skill or 2) driving the present value of the skill acquisition cost stream down will induce workers of lower marginal ability to invest in learning. The employer looking to hire skilled talent has both options, but using open source as a strategy offers an opportunity in the latter. In the short-run, TensorFlow drives the costs to learn how to train deep neural networks down, inducing marginally lower ability entrants to acquire the skill. Before addressing how firms benefit from this, we will solve for the firm's problem.

3.2 Firm Market Valuation with Intangibles

Here we follow a setup common to Lucas (1967); Hayashi (1982); Wildasin (1984); Hall (2001); Yang and Brynjolfsson (2001); Tambe et al. (2020a); Brynjolfsson et al. (2021). The formulation of the firm's market value problem is nearly identical to the presentation in Brynjolfsson et al. (2021)

[18] There are two special but relatively uninteresting cases: the first is if the skill is so difficult that nobody ever wants it (θ^* approaches ∞) and the second is if the skill is costless to acquire so everyone gets it ($\theta^* = 0$).

Appendix B and will follow that structure closely. We are just adding a variety of labor. Assume a production function of the following form:

$$Y = pG(A,\mathbf{K},L_H,L_L,\mathbf{I}) \tag{7}$$

Here Y is output, A is total factor productivity, p is the price of output, G is the production function, $\mathbf{K}$, L_H, L_L, and $\mathbf{I}$ are the vector of capital stock quantities by variety, the quantity of labor hired that chooses to invest in the skill, the quantity of labor hired that chooses not to invest in the skill, and the investment quantity vector by variety (respectively). Assume perfect competition between firms and constant returns to scale in all factor inputs. The vector of market prices for new investment goods is $\mathbf{z}$. G represents the final output net of adjustment costs, is non-increasing and convex in all varieties of $\mathbf{I}$, is non-decreasing and concave in all varieties of $\mathbf{K}$ and L, and has homogeneity of degree one. These restrictions make it such that the firm will pay opportunity costs at an increasing rate to accelerate capital investment, have diminishing marginal returns in single inputs, and constant returns to scale. The market value of a price-taking firm will be equal to the sum the assets of the firm, priced at the replacement cost plus the marginal adjustment cost per unit of assets. With j indexing capital varieties, the firm maximizes profits over investment in capital and hiring of perfectly flexible labor:

$$\max_{I,\ L} \left[\int_{t=0}^{\infty} \pi(t)\delta^t \, dt\right] = V(0) \tag{8}$$

where, suppressing subscripts for time t,

$$\begin{aligned} &\pi(t) = pG(A,\mathbf{K},L_H,L_L,\mathbf{I}) - w_H L_H - w_L L_L - \mathbf{z'I} \\ \text{and}& \\ &\frac{dK_j}{dt} = I_j - \beta_j K_j \ \forall j = 1,2,...,J. \end{aligned} \tag{9}$$

We add depreciation rates for capital variety j defined by β_j, and δ^t corresponds to the discount factor at time t as before. The firm solves for the solution to the Hamiltonian maximization at $t = 0$:

$$H(\mathbf{K}, L_H, L_L, \mathbf{I}, A) = (\pi(t)\delta^t) + \sum_{j=1}^{J} \lambda_j (I_j - \beta_j K_j)$$

with standard constraints (t subscripts are assumed and suppressed on all prices, quantities, and marginal products),

$$\begin{aligned}
&1)\ \frac{\partial H}{\partial \lambda_j} = \dot{K}_j = I_j - \beta_j K_j \ \forall j \in \ 1,2,...,J, \ \forall t \in [0, \infty), \\
&2)\ \frac{\partial H}{\partial K_j} = -\dot{\lambda}_j = pG_{K_j}\delta^t - \lambda_j \beta_j \ \forall j, \forall t, \\
&3)\ \frac{\partial H}{\partial I_j} = 0 = (pG_{I_j} - z_j)\delta^t + \lambda_j \ \forall j, \forall t, \\
&4)\ \frac{\partial H}{\partial L_H} = 0 = (pG_{L_H} - w_H)\delta^t \ \forall t, \\
&5)\ \frac{\partial H}{\partial L_L} = 0 = (pG_{L_L} - w_L)\delta^t \ \forall t, \\
&6)\ \lim_{t \to \infty} \lambda(\mathbf{t})\mathbf{K}(\mathbf{t}) = 0
\end{aligned} \tag{10}$$

Note here in constraints 4 and 5 above that labor earns precisely its marginal product and there are no adjustment costs for skilled labor. We will assume that ability is observable to both employer and employee, and that markets are competitive across all abilities. The two labor types are only separated out to illustrate that there are different varieties, but the human capital component of skilled worker compensation will require fixed costs of investment (it will be a capital variety, not labor). This leads to a value of the firm precisely equal to the sum of the individual capital variety values, priced at the shadow cost of investment:

$$V(0) = \sum_{j=1}^{J} \lambda_j(0) K_j(0) \tag{11}$$

The λ value for each capital variety includes its replacement cost as well as the present value of adjustment costs, priced at the marginal adjustment cost of competitors. For an asset like a cloud computing implementation, the replacement cost of the assets are observable in the market, but the market value of the installed cloud computing system within the firm includes the firm-specific component of the marginal product of that investment. Since competitors might not be able to generate a productive cloud system at as low a cost, the difference between what competitors can achieve and what the focal firm can achieve is firm-specific. Investors therefore price the asset above its replacement cost. λ prices on investment therefore have a component driven by the market price of investment z and a component that is the present value of the marginal product of the investment within the firm pG_I (suppressing subscripts).

This formulation is highly flexible. Nearly anything can be included as an asset in this framework, such as a monopsony rent capacity, complementarities between inputs, and intangible assets around business capabilities, branding, or managerial technologies. It also includes standard capital like property, plant, and equipment. Additionally, the model suggests that the marginal investment in a unit of capital will be equal to the competitors' marginal adjustment cost value (less replacement costs), but inframarginal rents are permitted.

3.3 Cheap(er) Learning and Firm Value

When workers acquire new technological human capital by investing in skills, they must rent their newfound human capital to their employer. At that point it becomes the employer's asset as long as the contract continues. We can then set the wages for labor of any type to be the same, and consider the residual premium awarded to employees with skills as a rent awarded to human capital. Assume then that $W_H = W_L = W$, but that there is a capital stock and investment variety K_H and I_H that the firm rents from the worker. Taking our case above, the compensation to workers who have purchased the new skill in period t is $W_t + z_{H,t}(\theta c_t)$ if each of these workers has θc_t units of I_H to supply at price $z_{H,t}$. The θ ability parameter converts costs spent to productivity. Denote efficiency units of θc_t then as $\tilde{c_{it}} = \theta_i c_{it}$ for each worker i.

Per efficiency unit, the firm sinks investment in human capital of its employees at a market investment price $z_{H,t}$ such that in each period the (discounted) difference between the marginal product of human capital investment and the market price is equal to the λ_H, the shadow price of human capital that includes the replacement cost and the firm-specific marginal adjustment costs of competitors. In other words, the employer invests in hiring employee human capital up until the point that its competitors can pay equal values for that human capital, the highest willingness to pay that the competitor business model facilitates. Where the employer can deploy its workers' skills in unique or hard to replicate ways, they will earn the excess above $z_{H,t}$ as a rent. Formally, from constraint 3 above:

$$\frac{\partial H}{\partial I_H} = 0 = (pG_{I_{H,t}} - z_{H,t})\delta^t + \lambda_H \ \forall t \tag{12}$$

For a firm hiring up to the M^{th} efficiency unit of I_H in period t, the total value to the firm of the employee human capital net of the costs paid for worker skill service flows from that period will be:

$$\int_i^M \tilde{c_{it}}[\lambda_H - \delta^t z_{H,t}]\, di = \int_i^M \tilde{c_{it}}[-p\delta^t G_{I_{H,t}}(i)]\, di \tag{13}$$

with the total value inclusive of what is paid to the worker as

$$V_{H,t} = \int_i^M \tilde{c_{it}}[\lambda_H]\, di = \lambda_H I_{H,t} \tag{14}$$

and straightforwardly, the total value to the firm of all investment in variety H is the discounted total $V_{H,t}$:

$$V_H = \int_0^\infty \delta^t V_{H,t}\, dt \tag{15}$$

The marginal product of investment $G_{I_{H,t}}$ is negative because investment constitutes foregone output in the current period, but is capitalized as future capital service flow value. The difference then between the shadow price λ_H and the discounted investment price paid to the employee for their skills constitutes the firm-specific rent (which will be equal to zero on the margin). On average, however, the per unit investment marginal product cost in foregone output is less than the shadow price value per unit. This generates the firm's Tobin's q value at a diminishing rate given our assumptions above.

Let us now consider the case that there is an open source launch of a tool like TensorFlow that makes it cheaper in present value terms for workers to pick up human capital.[19] From our equation for the ability level that is indifferent between acquiring and not acquiring the skill before, $\theta^*_{new} < \theta^*_{old}$, and therefore $F(\theta^*_{new}) < F(\theta^*_{old})$. With N total workers in the market, workers of lower ability will learn TensorFlow and supply in terms of workers who invest at all will increase from $N(1 - F(\theta^*_{old}))$ to $N(1 - F(\theta^*_{new}))$. Workers with some investment in ML skills will be incentivized to acquire even more from the cost shock. From the perspective of forward-looking firms, this will be stimulative for investment; rents paid to workers for skills will be lower at the new higher equilibrium quantity. But those investments, requiring foregone output, have not yet been made. That means this expected future increase in quantities of workers hired does not on its own increase market value according to the theory above.

Where might a market value effect arise? Since this change comes as a surprise to nearly all firms (except Google), the prior capitalized investment value prior V_H was made under assumptions of higher costs: $\mathbb{E}_{\text{before}}[z_{H,t}] > \mathbb{E}_{\text{after}}[z_{H,t}]\ \forall t$ if we set $t = 0$ at the time of the change. Therefore because of the change now, all prior investments are inframarginal, that is, $(pG_{I_{H,t}} - z_{H,t})\delta^t + \lambda_H > 0\ \forall t$ on the investments made before. Or, put differently, there is a price increase in the employer's human capital investment value because the ongoing costs to rent that capital are lower in perpetuity, but the expected future service flows are the same. **This is the core prediction of the theory: the launch of tools that make it easier to learn technological skills increases the firm market value via a price channel.** This formulation and possibility establishes the employer's sunk investments in human capital-intensive production processes as a form of real option on future employee skill sets. If the costs to acquire those skills are volatile (especially with a negative drift), the employer stands to benefit by refinancing their prior investment at lower human capital rental prices. This channel also suggests a difference-in-difference empirical specification. AI-intensive firms that have already sunk AI investments are exposed directly to the TensorFlow release, but firms without AI complements will not have yet to pay the fixed costs of investment and therefore have no opportunity to reprice their assets. With Δ to denote the difference in a measure between pre and post-open source launch and MV^{AI} and $MV^{\neg AI}$ denoting market value for firms with and without AI complements respectively, we might expect the following with no other effects or changes[20]:

[19] One might say this is deep learning becoming "cheap" learning.

[20] $I_H^{\neg AI}$ would be zero in this case.

$$\Delta MV^{AI} - \Delta MV^{\neg AI} = (\Delta z_H[I_H^{AI} - I_H^{\neg AI}]) \tag{16}$$

Other channels are also possible. The marginal product of investment might include exposure to A_t, the total factor productivity of the firm at time t. Increasing the firm's productivity has a similar effect of making the marginal product of I_H more negative since more output must be foregone to sink investments. Therefore all former investments made under lower productivity assumptions are more valuable. Importantly for this to affect market values on an ongoing basis, this must happen to all firms, but each in their own way. Otherwise the higher productivity of all assets will lead to higher market prices for investment, offsetting the marginal product effects. Either channel then relies on an argument for firm-specific assets, and likely intangible assets if they are difficult to replicate. Additionally, total factor productivity increases affect the value of all assets. This is empirically testable, as we will show with placebo skills and other assets.

It might also be the case that the new skill shock makes workers more productive inside their firms, driving up the efficiency units supplied by each ability level that has acquired the skill. This might have the opposite effect on asset prices depending on the firm demand elasticity. If, for example, one worker could supply all firms with her deep learning labor, then only the highest ability superstar worker would get paid for her human capital (and get paid a lot!) (Rosen, 1981).[21] Firms with sunk investments will possibly have done so assuming lower prices for skill "rental" than the market would bear. In this case the result is ambiguous. Lastly, it might be the case that there are dormant assets in complementarities between the firm's prior investments and the new skillset, were it more available (Choudhury et al., 2020). This kind of effect would entail a change in the production function itself without significant fixed costs. In a later section, I test this possibility by examining changes the market values of firms with high levels of worker exposure to ML, rather than high levels of skill.

To recap, the model suggests a few hypotheses that link cheaper employee skills to possibly increased firm market value. We might see 1) price channel effects as the employer "refinances" its human capital stock at a lower rate, 2) the new skill might increase total factor productivity at some point, causing a widespread asset price increase over all asset varieties with firm-specific components, 3) worker-level productivities might increase, driving up payments to skilled employees with ambiguous effects on market value but increased compensation to workers, and 4) business model changes could be facilitated by the cheaper capital variety. These possibilities are by no means unique to technology skills, but technology skills like deep learning are well-suited to an empirical study of possible channels. Tech skills value often depreciates quickly and open source software is linked to productivity (Deming and Noray, 2018; Horton and Tambe, 2019; Nagle, 2019). Additionally there is some evidence that IT-intensive firms carry greater asset risk (Dewan et al., 2007). Employee skills might account for some of that variation. We now turn to the data.

[21] This seems less likely given standard compensation rates for machine learning engineers. Superstars are still highly compensated, but there is a market for others.

4 Data on AI and Engineering Talent

Member profile information from LinkedIn serves as my main data source. Part of the reason the firm value of engineering and technological talent has been difficult to measure in the past is because of a relative paucity of granular data in this area. Online platforms like LinkedIn present an opportunity to tie organizations to the skills, education, career histories, and professional networks of their staff. Outside of governmental and administrative datasets, data at this scale and level of detail is unusual. LinkedIn has over 575 million members in over 200 countries and territories (more than 150 million U.S. members, 15 million in Canada, and 25 million in the U.K.). Additionally, over 26 million companies, 60 thousand schools, and 35 thousand skills are represented on LinkedIn.[22] The LinkedIn platform has become a standard tool for job seekers in many labor markets.

The primary data source for analyzing AI talent flows effects on firm market value come from LinkedIn-derived values from the relatively recently constructed panel of detailed skills data. LinkedIn first rolled out the skills product in 2011, though collection of high-fidelity records of member additions of skills began in 2013. LinkedIn now has over 55 million unique skills records across over 35 thousand standardized skill units. Recently, LinkedIn has categorized and standardized the approximately 35 thousand unique skills on its standard platform into a set of skills clusters using nonlinear embedding spaces.[23] These clusters are seeded and curated by human taxonomists and subsequently applied to co-occurrences of skills on profiles across the entire platform.[24] Skills are related by distance in "skill space" as a result of this machine learning-driven encoding. Skills that tend to be closer in this space are more likely to be associated together and tagged with a common human-curated cluster name. Likewise, skills that co-occur less frequently are classified in separate clusters. I make use of the production neural skills embeddings supplied by the LinkedIn engineering team. These embeddings are used actively by the company. For this analysis, LinkedIn fortunately had a number of pre-defined dedicated skill clusters related to AI, data science, and other relevant technology skills.

The result is a series of aggregated counts of skills additions in different categories which I then aggregate, accumulate, and normalize at the firm-year-occupation and firm-year levels. I extract specific skill counts for deep learning, machine learning, linear regression, and a handful of other data science skills. Skill records from individual profiles are aggregated by firm-quarter. Most specifications in the analysis log the accumulated skill counts from members (adding 1 to zero entries) because the skill counts are skewed (similarly to assets). All of these measures are then joined to Compustat measures of financial performance by firm and quarter. The Compustat measures serve

[22]Source: The LinkedIn Economic Graph Research and Insights team (https://engineering.linkedin.com/teams/data/data-science/economic-graph-research-insights). About 70% of platform membership is outside the U.S. The growth rate of membership is approximately 2 members per second as of July 2018. LinkedIn Economic Graph Research and Insights supplied data and valuable feedback to make this project possible.

[23]Clusters including Agronomy, Artificial Intelligence, People Management, and Digital Literacy (amongst others) and rely upon user-supplied data. Because the user-supplied data is highly variable, all skills go through a standardization algorithm before being made available for analysis.

[24]See https://en.wikipedia.org/wiki/Nonlinear_dimensionality_reduction for a set of useful embedding algorithms. TensorFlow can be used to build some of these models.

to create the primary outcome variables. Details on Compustat measure construction are in the Data Section of the Appendix. Figure 1 shows the aggregate skill additions for AI-related skills and advertising across the entire platform. There is some seasonality in the data, with more skills getting added in the beginning of the year. Table 1 shows some example skills for different aggregated categories. Table 2 reports summary statistics for skills by firm.

Figure 1: AI and Advertising Skill Additions Across LinkedIn by Month

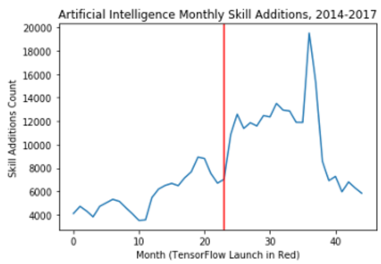

(a) AI Skills

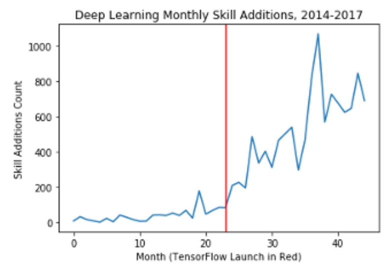

(b) Deep Learning

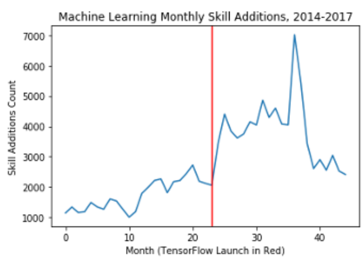

(c) Machine Learning

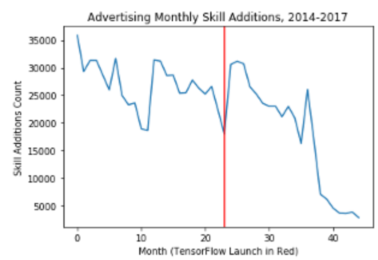

(d) Advertising

Figure Notes: These charts show the total additions of user-reported skills within publicly traded companies across the LinkedIn platform for the designated time period. The TensorFlow launch in red is November 2015.

Artificial Intelligence	Data Science	Digital Literacy	Advertising	Management	Data Storage	Cloud Computing
Artificial Intelligence	Forecasting	Excel	Advertising	Management	Microsoft SQL Server	Amazon Web Services
Machine Learning	Modeling	Word	Campaigns	Strategy	mysql	Microsoft Azure
Classification	Statistics	Powerpoint	Collateral	Strategic Business	sql	Google Cloud Platform
Information Retrieval	Analytics	Microsoft Office	Sponsorship	Small Business	Hadoop	
Computer Vision	Data Integrity	email	Direct Marketing	Strategic Planning	Databases	
Neural Networks	Statistical Tools	spreadsheets	Search Engine Marketing	Change Management	Datacenter	
Speech Recognition	Data Analysis	Windows	Brand Development	Executive Management	Storage	
Semantic Web	SAS	Mac	Media Planning	Service Providers	Data Warehousing	
Parsing	R	Lotus Notes	Email Marketing	Outsourcing	Hive	
Pattern Recognition	Sampling	Google Sheets	Media and Entertainment	Business Planning	Pig	

Table 1: Example Skills in Different Skill Clusters

Table Notes: Skill clusters taken from production embedding model and taxonomy output. Each skill cluster conceptually has overlap with others, but skills in this taxonomy are assigned specifically to one larger cluster. Deep Learning, for example, is part of the Artificial Intelligence cluster, but is not part of the Data Science cluster. To the extent that there is some misclassification of skills, there should not be measurable differences effects estimated for different skill indices — this facilitates the use of different indices as placebos in difference-in-difference analyses.

I also use the profile employment data. With over 180 million individual position records spanning from 2000 to 2017, I build firm-level aggregates of worker years of education as well as counts

Variable	Count (Firm-Quarters)	Mean	Std. Dev.	Min.	25th Perc.	Median	75th Perc.	Max
Value Added (Millions USD)	471508	220.63	1172.02	-71484.84	0.81	12.35	78.48	36657
Tobin's q Value (Millions USD)	530299	2017.88	12425.97	-190234.59	3.75	58.97	541.58	792807.63
Revenue (Millions USD)	473376	741.04	3723.4	-25623	3.86	38.59	245.91	207307.33
Total Human Capital (Years)	193426	101346.2	389103.67	1.33	2671.75	12283.67	58926.06	15149309
Wagebill-Weighted SML	194845	3.07	0.08	0	3.03	3.06	3.12	3.32
PP&E (Millions USD)	304094	2484.42	13894.4	-0.32	4.31	56.32	600.59	546691.81
Intangible Book (Millions USD)	442078	797.41	4893.85	-78.64	0	5.44	123.04	312052
Other Assets (Millions USD)	530299	6254.59	76651.92	-298653.78	1.25	50.62	540.39	3639525.8
Total Assets Book (Millions USD)	530299	9346.24	88752.59	-12.67	36.95	299.08	1651.4	3771199.8
Market Value (Millions USD)	530299	11364.12	91271.39	0	67.45	463.15	2481.19	3758335.5
Log(Market Value)	530299	6.07	2.58	-6.23	4.21	6.14	7.82	15.14
AB Testing Skill Count	32531	2.81	26.22	0	0	0	0	858
AI Skill Count	32531	102.57	1094.9	0	0	0	4	45930
Advertising Skill Count	32531	520.89	3697.22	0	0	12	108	150484
Big Data Skill Count	32531	48.73	546.41	0	0	0	2	26512
Cloud Computing Skill Count	32531	680.51	8697.12	0	0	3	43	341319
Data Science Skill Count	32531	1262.97	8546.92	0	3	31	259	326495
Data Storage Skill Count	32531	3574.41	37616.34	0	3	47	425	1549290
Deep Learning Skill Count	32531	1.46	24.79	0	0	0	0	1606
Digital Literacy Skill Count	32531	3314.98	18948.19	0	13	120	837.5	647816
Linear Regression Skill Count	32531	2.83	25.24	0	0	0	0	863
Management Skill Count	32531	5598.73	45191.02	0	20	159	1114	1771434
Log (AI Index)	32531	1.07	1.83	0	0	0	1.61	10.73
Log (Advertising Index)	32531	2.86	2.54	0	0	2.56	4.69	11.92
Log (Big Data Index)	32531	0.8	1.56	0	0	0	1.1	10.19
Log (Cloud Computing Index)	32531	2.17	2.43	0	0	1.39	3.78	12.74
Log (Data Science Index)	32531	3.63	2.69	0	1.39	3.47	5.56	12.7
Log (Data Storage Index)	32531	0.02	0.03	0	0	0.01	0.03	0.57
Log (Digital Literacy Index)	32531	4.7	2.88	0	2.64	4.8	6.73	13.38
Log (Management Index)	32531	5.04	2.86	0	3.04	5.08	7.02	14.39

Table 2: Summary Statistics by Firm

Table Notes: While all skill-related regressions use a completely balanced quarterly panel from 2014 to the end of 2017, the table reflects summary data for firms outside of the sample as well. For accounting data, the sample is larger. Value added is computed as the difference between revenues and cost of goods sold, residual market value is calculated as the market value less book value of assets. Many skill values at the firm level are 0, and skill additions tend to be highly skewed except for the most common skills like management and digital literacy. In this regard AI, Data Science, and other technology skills are very similar. The logged indices are actually $log(x+1)$ in all cases, where x is the skill count.

and total wage bills (employee counts multiplied by Bureau of Labor Statistics average wages) of specific varieties of worker. The process is similar to the variable construction in (Benzell et al., 2018b; Tambe et al., 2020a). The data appendix in Tambe et al. (2020a) describes the dataset in detail. The same data, in addition to skills data, are used in this paper.

The LinkedIn data covers a substantial portion of the global knowledge and human capital-intensive worker population. The representativeness of the LinkedIn panel is imperfect, with predictably sparser coverage of smaller (non-public) organizations, less educated workers, blue-collar workers, and non-U.S. firms. Further the sample quality varies by year as LinkedIn's adoption diffused through the workforce. While there are data going back substantially farther than 2000, the coverage at that point relies upon members populating their pages with highly detailed work histories. Additionally, the incentives governing whether to post certain information differ across workers. The selection of workers observed on LinkedIn is likely to differ in meaningful ways from the underlying employee population. Workers seeking employment, for example, are more likely to have updated employment history and skills information on their profiles.

With respect to skills data, incentives to report skill information are higher for job seekers. Another possible concern is whether or not the user self-reported data is accurate or reflects credential inflation. Especially for job seekers, dishonest reporting on LinkedIn is not costless, but it does occur. To the extent that my skill measures overstate the true skill levels within firms, the parameter estimates from the regressions to follow will represent lower bounds on the true effects. Perhaps more concerning, and more likely, is underreporting on the LinkedIn platform. It is difficult, especially with skills data, to construct a true estimate of the stock of corporate skills. Workers will commonly omit their qualifications on LinkedIn profile data. In this case, the regression estimates will constitute an upper bound; accordingly I interpret the skills measures not as precise measures of skill stocks, but rather as relative indices of these human capital varieties at the firm-level.[25]

I pursue a number of strategies to mitigate these potential sources of bias. The simplest is the inclusion of combinations of firm, industry-time, and time fixed effects in all regression specifications. In all specifications, however, I correct for occupation, year, and firm-based discrepancies between LinkedIn and administrative labor datasets from the Bureau of Labor Statistics Occupational Employment Survey (BLS-OES). The BLS-OES survey provides detailed industry-level measures of occupational employment and wage. As in Tambe et al. (2020a), I build a crosswalk between LinkedIn's internal occupational classification system and the BLS-OES Standard Occupational Classification (SOC) Code by year. While skills data is compiled into a firm-quarter index form, the overall human capital measures and explicit measures of talent varieties by firm are normalized using the BLS-OES survey in a procedure described in the Data section of the Appendix. Industry controls are especially important given the skewed distribution of technology skills across industries. Figure 2 below reports the distribution of AI skills by industry.

Occupations like software engineer, unsurprisingly, have high fidelity and near complete coverage

[25] Classical measurement error arising from underreporting, overreporting, and noise together will of course attenuate all estimated coefficients toward zero.

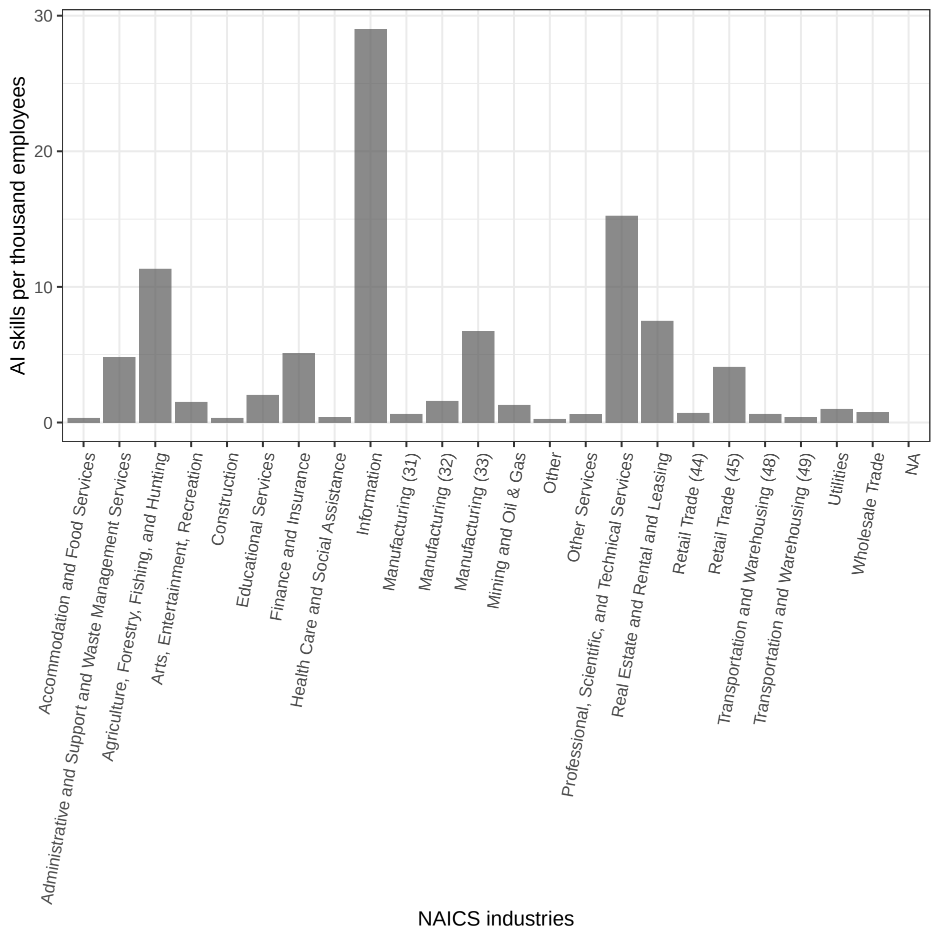

Figure 2: AI skills per thousand employees by 2-Digit NAICS

Figure Notes: This plots the prevalence of AI skills per thousand employees as of the end of the sample. The information, financial, advanced manufacturing, and scientific sectors have relatively higher penetration of AI skills. Accommodation and food services, construction, and transportation and warehousing have relatively fewer AI employees.

for U.S. firms. A few other titles, like dentist or transportation specialists, have lower baseline levels of coverage but are adjusted to BLS-OES consistency with this process. Nevertheless, the occupations and firms for which LinkedIn membership is relatively sparse will have noisier adjusted employment shares as well. To handle these issues as well as to effectively tackle the research questions in the paper, fixed effects at the industry, year, and firm-level are included in regression specifications. For engineering, research, and IT worker stocks, the relative presence on LinkedIn is higher in comparison to other occupations.[26] Software engineers and related occupations constitute the vast majority of workers with AI skills. But for analyses of other skillsets these sampling issues are potentially more concerning. In any case, the residual variation in coverage issues that might bias estimation coefficients for models with firm and year effects, for example, must change within firm and over time (and analogously within industry-year for specifications with those fixed effects). Plots and statistics on LinkedIn's coverage are reported in Tambe et al. (2020a) as well as the Data section of the Appendix.

Given the theoretical emphasis on sharp changes in a specific type of human capital, variation in firm-level general firm-level human capital stocks is a likely confound. I construct a total education years variable for each firm in each year as a control for the overall level of human capital at each firm. For this variable, following Tambe et al. (2020a), I aggregate the educational records of the workers according to the years of education required to achieve each listed degree.[27] That is, an Associate's degree counts as two years, a Bachelor's degree counts for four years, a Master's degree counts for two years, a research doctorate or medical doctor degree counts as six years. High school, for an alternative measure of education years, is counted as 12 years. These values are adjusted for coverage in the procedure above, and summed by firm-year to generate a total education years control.[28] Descriptive statistics for the LinkedIn measures can be found below in Table 2. One of the possible outcomes of the TensorFlow shock is that companies with high exposure to machine learning in their workforces would be affected. This ML exposure might be driven by automation and/or augmentation opportunities. I use the Suitability for Machine Learning (SML) measure from (Brynjolfsson et al., 2018b, 2019) to test whether TensorFlow created or destroyed market value in companies that employ workers with machine learning exposure. This measure is constructed at the 6-digit SOC code occupational level and reflects aggregated machine learning scores from 1 (low exposure) to 5 (high exposure) for a rubric with about 20 questions. The measure is not designed to represent automation, but rather the relative exposure of occupations to ML technologies. With measures of counts of employees in different occupational categories by firm and quarter, I create weighted-average SML scores for each firm in each quarter using BLS-OES salaries to construct wagebill (employment count times annual salary) weights at the firm-quarter level. SML is forward-looking in the sense that it represents potential. It is therefore useful in testing the hypothesis

[26]IT, Research, and Engineering are all defined as functional areas by LinkedIn. Workers of specific job titles in specific job functions are mapped into these functional areas. I aggregate counts, wage bills, and human capital after applying the normalization procedure detailed above.

[27]The normalization of education years to adjust for coverage is an identical process to the count normalization process described in the data appendix.

[28]This can be considered as the answer to the question "how many years has the firm gone to school?".

that firms with a highly specific form of intangible complement to AI — workforce reorganization potential due to ML — are better positioned to benefit from improvements in AI skills technologies like TensorFlow and PyTorch.

5 Estimating the Market Value Shock from AI Talent

5.1 Hedonic Market Value Analysis

Equation 11 above suggests a natural decomposition of the market value of firms onto the valuations of their constituent assets. This is the first step to determining the contribution of AI to firm's Tobin's q value. But a simple predictive OLS specification is useful for another purpose: if AI talent is not correlated with market value or other measures of firm performance, then any following analysis is likely unnecessary. Tables 3 and 4 below detail some suggestive evidence that it is indeed the case that AI and other skills are correlated with improved firm performance. In Table 3, various classes of the following regression model are estimated:

$$MV_{it} = \beta_{TA} TA_{it} + \beta_{HK} HK_{it} + \beta_{AI} AI_{it} + \mathbf{X'_{it}} \boldsymbol{\gamma} + \mu_i + \nu_t + \epsilon_{it} \tag{17}$$

The included variables in equation 17 are TA, the total book value of assets at the firm, HK, the (logged) human capital index built from adding up the total years of education for all workers whose LinkedIn profiles indicate being employed at firm i at time t (normalized following the procedure in the appendix and Tambe et al. (2020a)), the AI skills index measure, a firm fixed effect, a time fixed effect, and X, a matrix of additional skills indices and controls[29]. All skills indices are constructed as the cumulative sum of skills reported by workers of that variety in that firm-quarter. These indices are logged when nonzero (1 is added to entries with zero). In Table 4, model specifications with firm and industry-time fixed effects are compared for outcome variables of market value, revenue, and value added (respectively). That leads to a nearly identical setup as equation 17 but for different outcomes of Revenue and value added in specifications 3-6. These additional outcome variables are for the purposes of examining possible correlations between AI and firm performance. Logged outcome specifications will be additionally useful for testing whether or not the skills indices are related to the total factor productivity residual.

Using a balanced panel of publicly traded firms for all quarters from 2014-2017 and a variety of skills indices, a regression of market value on these logged skills indices returns a coefficient of \$480.7 million (standard error \$188.7 million) per 100% increase in the AI skills index for specification (1) with firm and year fixed effects. Other point estimates in the other specifications vary between approximately \$400 million and \$410 million with similar precision, and the specifications vary by introducing additional skills indices as controls. These estimates for AI are all significant at a 5 percent level. The one exception is if previous market value is used as a control. In this case the point estimate is positive and large at \$185 million but no longer statistically significant. This more

[29] As well as a vector of 1s for the intercept

Hedonic Regressions (Firm FEs)	(1) Market Value	(2) Market Value	(3) Market Value	(4) Market Value	(5) Market Value	(6) Market Value
Total Assets	1.080***	1.079***	1.079***	1.079***	1.079***	0.870***
	(0.0472)	(0.0472)	(0.0472)	(0.0472)	(0.0472)	(0.0498)
Log(Human Capital)	1,512***	1,483***	1,485***	1,488***	1,487***	1,068***
	(560.1)	(548.6)	(548.4)	(548.6)	(543.5)	(346.8)
Log(AI Index)	480.7**	398.8**	407.2**	409.0**	406.2**	185.0
	(188.7)	(190.7)	(189.5)	(189.2)	(190.4)	(169.0)
Log(Data Science Index)		567.1*	579.1**	585.1**	462.2**	309.9*
		(293.1)	(290.0)	(290.5)	(227.0)	(176.0)
Log(Cloud Computing Index)			-51.93	-47.74	-87.05	-94.49
			(160.3)	(160.7)	(157.6)	(123.2)
Log(Data Storage Index)				-4,766	-4,378	-641.2
				(7,383)	(7,389)	(5,706)
Log(Digital Literacy Index)					-3.068	18.28
					(193.2)	(160.7)
Log(Management Index)					288.1	261.1
					(359.8)	(318.1)
Log(Advertising Index)					141.2	104.3
					(123.6)	(98.17)
Lagged MV						0.259***
						(0.0657)
Constant	-10,658*	-12,484*	-12,445*	-12,406*	-13,795*	-11,816**
	(6,230)	(6,932)	(6,940)	(6,937)	(8,352)	(5,744)
Observations	23,944	23,944	23,944	23,944	23,944	23,904
Firm and Quarter FE	✓	✓	✓	✓	✓	✓
Industry-Year FE	X	X	X	X	X	X

Table 3: Market Value Regressions on Skill Indices

Table Notes: *** p<0.01, ** p<0.05, * p<0.1. Industry is defined at the 4-digit NAICS level. Standard errors are clustered by firm for all specifications. The table reports firm fixed effect ordinary least squares (OLS) regressions of market value on skill indices and associated covariates (total assets, log(total years of education in the firm) as human capital, and lagged market value. The logged indices are actually $log(x + 1)$ in all cases.

Firm Performance	(1) Market Value	(2) Market Value	(3) Revenue	(4) Revenue	(5) Value Added	(6) Value Added
Total Assets	1.079***	1.006***	0.0340**	0.0132***	0.0123**	0.00815***
	(0.0472)	(0.0206)	(0.0132)	(0.00421)	(0.00506)	(0.000934)
Log(Human Capital)	1,487***	2,440***	177.0***	868.9***	79.43***	245.2***
	(543.5)	(469.6)	(49.36)	(152.4)	(21.59)	(34.53)
Log(AI Index)	406.2**	4,985***	-1.250	691.6***	19.28*	262.0***
	(190.4)	(892.1)	(41.04)	(165.0)	(10.19)	(48.72)
Log(Data Science Index)	462.2**	-594.1	18.89	-333.5*	2.314	-142.6***
	(227.0)	(588.9)	(17.66)	(176.7)	(7.622)	(52.34)
Log(Cloud Computing Index)	-87.05	779.7	-17.76	176.3	-6.378	93.94**
	(157.6)	(541.3)	(42.99)	(110.5)	(9.558)	(38.05)
Log(Data Storage Index)	-4,378	-68,537***	-973.3	-9,725**	727.3	-3,992**
	(7,389)	(25,845)	(1,422)	(4,393)	(460.0)	(1,572)
Log(Digital Literacy Index)	-3.068	394.8	18.75	-77.87	2.798	-17.43
	(193.2)	(1,134)	(27.42)	(121.8)	(9.164)	(50.28)
Log(Management Index)	288.1	-1,783*	3.004	-57.27	31.47*	-26.15
	(359.8)	(1,039)	(84.59)	(168.3)	(16.38)	(60.41)
Log(Advertising Index)	141.2	1,920**	32.36	203.2	4.971	108.7**
	(123.6)	(751.8)	(47.02)	(146.5)	(6.731)	(41.85)
Constant	-13,795*	-19,979***	-826.3	-6,931***	-651.8**	-2,040***
	(8,352)	(4,918)	(680.5)	(1,353)	(327.6)	(362.5)
Observations	23,944	23,196	21,885	21,122	21,855	21,090
Firm and Quarter FE	✓	X	✓	X	✓	X
Industry-Year FE	X	✓	X	✓	X	✓

Table 4: Firm Performance Regressions on Skill Indices

Table Notes: *** p<0.01, ** p<0.05, * p<0.1. Industry is defined at the 4-digit NAICS level. Standard errors clustered by firm for firm fixed effects specifications and clustered by 4-Digit NAICS code for industry-time specifications in parentheses.

valuable firms are more likely to invest in AI, however, then lagged market value is a possible bad control (though still useful to include for comparison's sake). Nevertheless, if market prices are perfectly efficient, this table reflects a pricing of AI talent at the firm level. All of these estimates include firm and quarter fixed effects, so the variation is within-firm.

The other results in Table 3 are encouraging for validation purposes. Coefficients on Total Assets are roughly equal to $1 per dollar on the balance sheet, the human capital index coefficients are positive, large, and significant, and many of the skillsets that are relatively abundant like digital literacy and cloud computing have no statistically significant relationship with market value after including firm fixed effects. Comparing the full specification of skills indices with different types of fixed effects is informative, as the firm performance measures in Table 4 do. Column (1) in this table is the same as column (5) in the first table. Notably we see for market value, the industry-time fixed effects (industry defined as 4-Digit NAICS code) specification in column (2) places an even higher valuation on AI skills. In this specification, a 1 percent increase in the AI skillset of a firm is correlated with a market value increase of about $50 million.

The reason the firm fixed effects make such a large impact on the coefficient estimates is predicted by q theory of investment. With industry-time fixed effects, the specification is closer to matching the valuation maximization problem of the firm. In the theory, firms invest up until the point that the additional q value they create is equal to the marginal adjustment costs of competitors (what the firm can do that the competitors cannot with the same assets). Therefore controlling for time varying industry proxies for the best capabilities of competitors to dissipate rents. Including firm fixed effects, on the other hand, forces any covariance between the firm's characteristics and the performance outcomes to arise from changes within the firm — any fixed assets driving the results will be stripped away from the AI skill index. Firm fixed effects are therefore a possibly overly conservative (even adversarial) modeling choice to finding rents. They remove the time-invariant asset base and econometrically make the firm compete with itself for q value in the residual time-varying investment. Comparing specifications across different types of fixed effects models is therefore a useful approach.

In specifications (3) and (4), the predicted output is revenue instead. A 1 percent increase in the AI skills index is correlated with an approximate $6.92 million (standard error $0.17 million) increase in revenues for specification (4) with industry-time effects, but including firm fixed effects reduces this point estimate considerably and leads to a loss of statistical significance. Value added in the remaining columns shows consistent results with the other outcome variables, though the coefficient on AI skills for the firm effect model remains significant at a 10 percent level. Again, the variation in firms drives the variation in revenues and adoption of AI. Firms with greater revenues tend to invest more in AI talent, adjusting for firm size.

5.2 Main Difference-in-Difference Results: AI Skills Increase Market Value

If learning how to do deep learning specifically, and machine learning more generally, has really become substantially cheaper, does this skills acquisition cost reduction show up in the market value of employer firms? The previous section established a correlation between AI skills and market value adjusting for a wide series of technology and other related skills. AI appears correlated with valuations in a statistically and economically significant way. The next step is to apply a series of difference-in-difference approaches to evaluate the estimate for a causal impact of the increasing abundance of AI skills in the labor market. If we consider the launch of TensorFlow as a natural experiment, we have the following difference-in-difference setup following the theoretical prediction made in equation 16. For the empirical estimation, however, we include a series of like confounds via the other skills indices to isolate the AI investment effect. In particular,

$$MV_{it} = \beta_1 TA_{it} + \beta_2 HK_{it} + \beta_3 AI_{it} + \beta_4 [PostTensorFlow_t * AI_{it}] + \mathbf{X'_{it}}\boldsymbol{\gamma} + \mu_i + \nu_t + \epsilon_{it} \qquad (18)$$

This is a very similar regression to that in equation 17, but equation 18 modifies the approach to an ordinary two-way fixed effects difference-in-difference estimation. The coefficient of interest for the causal effect of the TensorFlow shock, if all assumptions necessary for identification hold, is β_4. That will reveal the effect of AI talent in the post period. Firm and time unit fixed effects are also included (μ_i and ν_t) to adjust for unit-specific and time-specific differences in means across firms. This limits the kind of variation that could lead to market value effects to be within-firm; cross-sectional differences in firms that predict Tobin's q ideally should not affect this estimate (I will discuss possible challenges to identification in a later section). The estimate in this equation will describe the additional market value increase to firms with assets that are complementary to AI in the post period. After adjusting for time-varying AI investment and other skills, firm size in terms of total assets, firm size in terms of overall human capital accumulation[30], and other controls, an increase in the market value of firms with AI complements relative to those without in the post period provides evidence of a causal effect of AI skills becoming easier to acquire. The assumptions necessary for a conclusion of this sort are that 1) AI skills are indeed becoming more abundant and easier to acquire in the post period, 2) that the trends of corporate valuations for firms with and without AI complements before AI skills became easier to acquire are parallel (i.e. firms that have complementary assets would be worth less in a counterfactual world without TensorFlow), 3) conditional on the included covariates (including AI and other skill indices), the exposure to TensorFlow was not determined by market value (as might be the case if some high market value firms got early access), and 4) the stable unit treatment value assumption (SUTVA). I will revisit these assumptions, but for now will note that Google is dropped from all analyses to avoid an obvious possible cause of endogeneity. The results are below in Table 5.

Table 5 reports economically and statistically significant impact for effects of AI talent in the post period. The coefficients on the AI Skills Index x TensorFlow post-period variable vary between \$751 million (standard error \$268 million) in column (6) to \$1.12 billion (standard error \$438 million) in column (5). Each column varies only by the included control variables, though all specifications

[30]This adjustment is made via the logged total education years at the firm, HK_{it}.

AI Skills Difference-in-Difference	(1) Market Value	(2) Market Value	(3) Market Value	(4) Market Value	(5) Market Value	(6) Market Value	(7) Market Value
Total Assets	1.089***	1.089***	1.090***	1.090***	1.090***	0.878***	0.876***
	(0.0545)	(0.0545)	(0.0545)	(0.0545)	(0.0544)	(0.0533)	(0.0538)
Log(Human Capital)	1,602***	1,594***	1,613***	1,617***	1,640***	1,151***	1,153***
	(587.8)	(583.1)	(587.0)	(587.5)	(588.6)	(381.1)	(381.5)
Log(AI Index)	-1,049*	-1,070*	-1,017*	-1,014*	-1,007*	-816.7**	-822.2**
	(593.0)	(600.0)	(588.9)	(588.2)	(584.9)	(398.4)	(398.6)
Log(AI Index)xPostTF	1,106**	1,101**	1,111**	1,112**	1,118**	751.1***	755.5***
	(434.7)	(433.5)	(435.5)	(435.6)	(437.8)	(268.2)	(268.2)
Log(Data Science Index)		186.3	278.1	287.0	183.3	122.7	111.7
		(229.5)	(236.1)	(236.0)	(215.5)	(178.2)	(179.6)
Log(Cloud Computing Index)			-403.1**	-394.4**	-436.7**	-345.4**	-345.8**
			(190.4)	(189.9)	(200.1)	(143.1)	(143.5)
Log(Data Storage Index)				-9,203	-9,093	-4,123	-4,135
				(7,700)	(7,773)	(5,766)	(5,787)
Log(Digital Literacy Index)					-228.3	-127.9	-123.2
					(179.4)	(166.0)	(165.9)
Log(Management Index)					681.4	548.2	534.1
					(470.4)	(398.7)	(396.8)
Log(Advertising Index)					4.805	-5.366	0.575
					(140.9)	(114.5)	(114.3)
Lagged MV						0.270***	0.272***
						(0.0662)	(0.0670)
Lagged MV Growth							-809.5*
							(419.1)
2 Qtr Lagged MV Growth							-295.3
							(291.3)
Constant	-10,562*	-11,191*	-10,863*	-10,779*	-13,039	-11,608**	-11,537**
	(6,173)	(6,533)	(6,464)	(6,456)	(8,122)	(5,614)	(5,570)
Observations	21,275	21,275	21,275	21,275	21,275	21,273	21,260
Firm and Quarter FE	✓	✓	✓	✓	✓	✓	✓

Table 5: Main Results: TensorFlow Difference-in-Difference

Table Notes: *** p<0.01, ** p<0.05, * p<0.1. Standard errors clustered by firm in parentheses. The table reports results from estimating versions of equation 18 on a balanced panel of firms from the beginning of 2014 to the end of 2017. The logged indices, including human capital, are actually $log(x + 1)$ in all cases. Column (1) reports a specification with no additional adjustments other than the fixed effects, total assets, and human capital. Columns (2) to (5) add in different skills indices, including advertising (in case AI is mostly useful for ad tech). Columns (6) and (7) include controls for lagged market value (1 quarter) and lagged market value growth (1 and 2 quarters) respectively in order to further control for possible trends in market value like AI hype or momentum in investment. Specifications including unit fixed effects and lagged dependent variables are known to create downward bias in coefficient estimates. These biases are adversarial finding a large effect, but qualitatively results are preserved in these specifications. Overall the AI index x Post-TF estimates are relatively stable between $7.5 and $11.2 million more market value per 1 percent increase in AI skills in the post period. In other words, TensorFlow appears to cause proportionate increase of $750 million to $1.1 billion per 100 percent increase of AI talent in companies. Unbalanced panel and Tobin's q (market value less book value of assets) results are reported in Tables 13 and 14 in the Appendix.

include adjustments for the total book value of assets in the firm, the logged total education years at the firm variable (human capital), and firm and quarter fixed effects. While the theoretical coefficient on total assets without adjustment costs is $1 per unit of asset book value, in reality this can vary somewhat. The coefficients on total assets are not statistically significantly different from this theoretical value in columns (1) to (5). But the inclusion of total assets enforces that the other covariates are used to explain variation in the Tobin's q value of the firm. Columns (2) through (5) each add an additional skill index control. Columns (6) and (7) also include controls for one quarter of lagged market value (6) and two quarters of lagged market value growth (7). The lagged market value covariate estimate is statistically significant at a 5 percent level. These covariates are designed to partial out the effects of momentum or AI-related hype. If AI companies are simply becoming more valuable, then the simple effect of being a high return company is removed with these effects. Specifications including unit fixed effects and lagged dependent variables are known to create downward bias in coefficient estimates. These biases are adversarial finding a large effect, but qualitatively results are nevertheless preserved in these specifications. Coefficients on the AI index itself are negative and almost perfectly offset by the post period coefficient, suggesting the most of the returns to AI investment have occurred following the proliferation of workers with AI talent via TensorFlow and other packages.[31] This set of results establishes that there is a post-period effect, though does not assign a mechanism. Indeed, the effect might not even coincide with the TensorFlow launch. I estimate therefore a new regression, interacting the AI skills index with each time period to create an event study version of equation 18:

$$MV_{it} = \beta_1 TA_{it} + \beta_2 HK_{it} + \beta_3 AI_{it} + \sum_{t=1}^{15} [\beta_{4t} AI_{it} * Z_t] + \mathbf{X'_{it}}\gamma + \mu_i + \nu_t + \xi_{it} \tag{19}$$

In this equation, each β_{4t} is unique for the time period t. The coefficient values are reported in the Appendix in Table 11. For event study specifications, I include total assets, human capital, and the full set of other skills indices related to technology following column (5) of Table 5. The results are summarized in Figure 3 below:

While adding this many additional covariates on the right-hand side of the regression increases the standard errors, the point estimates are consistent with the timing of market value increases coinciding with the TensorFlow talent shock. Market Values for AI using firms relative to the 0 benchmark of firms that do not use AI at all (AI Index measures multiplied by time dummies for firms with zero AI will always be zero). It is clear from the figure that companies with assets complementary to AI are growing in value relative to the companies without them, and that this growth begins sometime around the period in which AI talent began to proliferate towards the end of 2015. There are do not seem to be notable trends prior to the TensorFlow launch quarter. It is worth noting that since market value is a forward-looking measure, firms with AI-complementary assets need not be able to find immediately available talent to benefit from the TensorFlow launch.

[31] The offset of the AI Index effect in the post-period is an indication that the there is an AI-related increase in value following the TensorFlow launch, and that AI intangibles are negatively correlated with valuations in the pre-period.

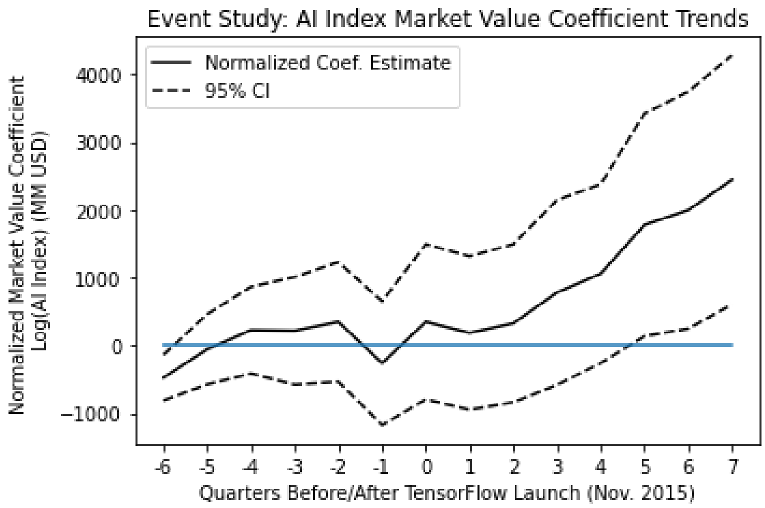

Figure 3: Event Study Results - Full Balanced Sample

Figure Notes: This plots the 95 percent confidence interval for the coefficients β_{4t} from equation 19 estimated on the set of covariates in column (5) of Table 5 over time. All values have been normalized by subtracting the pre-TensorFlow mean value of the coefficients to set an (average) baseline value of 0 in the pre-treatment period. The zero benchmark represents firms that do not use AI at all. The upper and lower bounds of the 95 percent confidence interval are represented with dashed lines, and the 0 line is in blue. The y-axis is the coefficient values for the regression of market value on the interacted time dummies and AI skills. A coefficient of 1000, for example, means that a 100 percent increase in AI in that period predicts a $1 billion increase in market value. Standard errors are clustered by firm.

Instead what is necessary is for forward-looking investors to reprice the firm's present investments on the assumption that in the future talent will be less of a bottleneck. What this figure and the foregoing analysis does not reveal is which of the hypothesized mechanisms leads to the market value change. The next section discusses and investigates those possibilities.

6 Testing Possible Mechanisms for the AI-Related Market Value Increase

6.1 Price, Productivity, and Corporate Workforce ML Exposure Results

Revisiting the hypotheses mentioned in Section 3, there are a number of possible mechanisms via which a change in the difficulty of acquiring a new technical skill for employees might affect the market value of employers. The first possibility is price. Forward-looking investors judge that installed asset base of the firm is worth more after one of the necessary complementary inputs (worker human capital) becomes more abundant. Since the expected future rental rate of those skills is lower if more workers have them, the firm's investments sunk under the assumption that the rental rate would be higher going forward are now a source of quasi-rents. New investment, however, will require paying the market rate for additional capital (human and otherwise), which will be higher. The second possibility is an increase in total factor productivity.

By a similar logic to the price shift, if firms are now more productive as a result of their employees' newfound ability to apply a technical skill, then the installed asset base will yield more in terms of capital service flows than it was expected to do. Two types of productivity change should be separated. An instantaneous productivity increase from, for example, existing deep learning workers shifting their approach to a more efficient platform, increases demand for complementary skills like data science or cloud computing. If productivity is growing and causing a demand-side effect, then the valuation of skills other than deep learning (but perhaps in a similar category) should also grow. An expected future productivity increase, on the other hand, falls under the price mechanism category as well, but affects all asset varieties with firm-specific components. Without any firm-specific component, competitors will be able to bid up the value for assets that would be equally productive in any company. The firm-specific complementary assets to AI talent in this case grow more valuable because of what the firm itself is capable of doing. It might also be the case that worker-level productivity changes. This is a likely result of a new software package or easier method of approaching complicated problems. The net effects of such a change on market value are ambiguous, as more productive workers can achieve more for their employers but also demand more compensation. Without detailed compensation data, it is difficult to document the extent to which worker productivity is enhanced, but Horton and Tambe (2019) provides some useful discussion of related topics.

Lastly, it might be the case that the business model, including intangible assets like culture, business processes and procedures, intellectual property, customer relationships, and talent strategy is exposed to AI in a different way. For instance, a competitor discovering a new means of competing with fully automated production processes poses a risk to a company with many employees in the same market. This has occurred in the past, such as when eBay launched automated translation with ML tools and increased international trade on its platform (Brynjolfsson et al., 2018a). Often many of these business-specific assets are likely to be time-invariant and firm fixed effects remove them from the foregoing analysis. But sensitivity to machine learning technology in the workforce by firm is measurable (Brynjolfsson and Mitchell, 2017; Brynjolfsson et al., 2018b; Felten et al., 2018; Brynjolfsson et al., 2019; Webb, 2019) and varies over time with the composition of the workforce. Particularly for the case of AI/ML, increasing the workforce using the technology might have implications for firms and workers that are exposed to the technology in other ways (Acemoglu et al., 2020). I apply the SML measure in Brynjolfsson and Mitchell (2017); Brynjolfsson et al. (2018b, 2019) to discern whether it is exposure to AI talent or talent potentially affected by AI that leads to market value changes in the post period.

Evidence for the productivity hypothesis in other research is thin. This is to be expected, as the productivity benefits from new technologies, especially significantly impactful technologies, can take years to show up (Brynjolfsson et al., 2021). Babina et al. (2020) and Acemoglu et al. (2020) find little correlation between firm productivity and AI talent using alternative datasets, including job postings. My results here using data on stocks of skills (as opposed to evidence of AI demand from postings) is in agreement with these other studies.

Table 6 below reports the result of regressions of logged revenue and logged value as outcomes on logged skill indices and the asset and human capital controls. These are standard ordinary least squares productivity regressions. If we assume a Cobb-Douglas production function, then $Y = AK^{\alpha}L^{1-\alpha}$ where K and L are capital and labor (in this case proxied by total assets and the human capital measure since that is likely highly correlated with labor inputs). A is the total factory productivity, and in a log-log regression we have:

$$y_{it} = \beta_0 + \beta_1 k_{it} + \beta_2 l_{it} + a_i t + \eta_{it} \tag{20}$$

Lower case variables in equation 20 denote logged variables. The residual of a regression of Y (either revenue or value added) on logged capital and labor will contain productivity. Therefore if AI skills (or other skills) are statistically significant when estimating equation 20, they might be correlated with productivity A. As Table 6 shows, there is little evidence of AI's correlation with contemporaneous productivity in the 2014 to end of 2017 sample period.

Total assets and human capital are both strongly correlated with the outcome variables, but there are no uniformly present correlations between the skill indices selected and the outcomes. For firm fixed effect models, this may suggest that productivity effects of these technology skills in the sample period are absorbed by fixed firm characteristics. The industry-time model point estimates are somewhat more positive. But overall it does not seem that AI talent in particular is currently correlated with productivity. None of the coefficients on AI skills are statistically significant. Further investigating the productivity hypothesis will be an analysis of placebo skills in Section 7.

To test the specific hypothesis that the relevant intangible asset is workforce exposure to future AI opportunities, I apply a similar difference-in-difference analysis as in the main results following equation 18 testing the difference-in-difference estimates for SML instead of AI skills and equation 19 conducting an event study on the interaction between time dummies for each period in the sample and the firm-level SML score. Firm-level SML scores are calculated in a multi-step process. First I calculate the wagebill paid by each firm to each occupation in each quarter using a normalized count of LinkedIn profiles at each firm in each quarter by each occupation (defined by 6-digit SOC Code), multiplied by the annual salary for that occupation reported by the BLS-OES (this is the procedure described in the data appendix of this paper and of Tambe et al. (2020a)). These wagebills by occupation-quarter constitute the weights in a weighted average of occupation-level SML scores (also for 6-Digit SOC Codes). What results is a wagebill-weighted overall SML score for the firm that reflects the central tendency of the SML scores given how much the firm spends on each type of worker. The results of a difference-in-difference estimation procedure with SML score instead of AI skills are reported in Table 7 below.

Coefficient estimates for the SML and SMLxPostTensorFlow covariates are not statistically significant for any of the specifications, and the point estimates are negative for all specifications except for column (6). Given that the requisite assumptions hold for identifying a causal effect of AI exposure, these results are evidence that so far ML exposure in the workforce is not a channel via which TensorFlow might have caused an increase in employer market value. Bolstering the case that high SML firms are not affected positively by the TensorFlow change (or other AI talent

Productivity Regressions (OLS)	(1) Log(Revenue)	(2) Log(Revenue)	(3) Log(Value Added)	(4) Log(Value Added)
Log(Total Assets)	0.544***	0.661***	0.581***	0.731***
	(0.0336)	(0.0315)	(0.0360)	(0.0259)
Log(Human Capital)	0.0986***	0.313***	0.0863***	0.198***
	(0.0176)	(0.0290)	(0.0182)	(0.0237)
Log(AI Index)	-0.00312	-0.0170	0.00275	-0.00210
	(0.00581)	(0.0130)	(0.00691)	(0.0122)
Log(Data Science Index)	0.0165	-0.0165	0.00651	0.0293
	(0.0166)	(0.0260)	(0.0187)	(0.0211)
Log(Cloud Computing Index)	-0.0276***	0.0112	-0.0302**	0.0425**
	(0.00997)	(0.0208)	(0.0118)	(0.0177)
Log(Data Storage Index)	1.502*	-0.318	1.234*	0.488
	(0.787)	(0.777)	(0.630)	(0.862)
Log(Digital Literacy Index)	0.0157	0.0541**	0.0202	0.0298
	(0.0173)	(0.0225)	(0.0203)	(0.0241)
Log(Management Index)	0.0286	0.00770	-0.00304	-0.0547**
	(0.0258)	(0.0393)	(0.0311)	(0.0265)
Log(Advertising Index)	0.0122	-0.0305***	0.0131	9.17e-05
	(0.0151)	(0.0115)	(0.0129)	(0.0159)
Constant	0.346	-2.617***	-0.597*	-2.900***
	(0.298)	(0.229)	(0.306)	(0.118)
Observations	21,704	20,938	21,007	20,247
Firm and Quarter FE	✓	X	✓	X
Industry-Year FE	X	✓	X	✓

Table 6: OLS Productivity Regressions

Table Notes: *** p<0.01, ** p<0.05, * p<0.1. Standard errors clustered by firm for firm fixed effect specifications and by 4-digit NAICS for industry-time fixed effect specifications in parentheses. The table reports results from estimating versions of equation 20 on a balanced panel of firms from the beginning of 2014 to the end of 2017. The logged indices, including human capital, are actually $log(x+1)$ in all cases. Columns (1) and (2) reports the elasticity of revenue with respect to total assets, human capital, and skills indices. Columns (3) and (4) report the same for value added.

SML Difference-in-Difference	(1) Market Value	(2) Market Value	(3) Market Value	(4) Market Value	(5) Market Value	(6) Market Value
Total Assets	1.081***	1.081***	1.080***	1.080***	1.080***	0.871***
	(0.0475)	(0.0474)	(0.0474)	(0.0474)	(0.0474)	(0.0498)
Log(Human Capital)	1,537***	1,487***	1,481***	1,483***	1,475***	1,036***
	(553.1)	(538.7)	(538.3)	(538.6)	(530.7)	(332.6)
SML (Wagebill-Weighted)	-917.7	-626.3	-560.5	-553.3	-614.4	1,139
	(6,937)	(6,841)	(6,840)	(6,843)	(6,798)	(5,315)
SMLxPostTF	-1,946	-1,809	-1,823	-1,833	-2,024	-2,411
	(2,080)	(2,032)	(2,031)	(2,030)	(2,069)	(1,657)
Log(Data Science Index)		746.7**	715.8**	721.0**	564.5***	339.1**
		(291.2)	(283.4)	(283.9)	(212.4)	(162.8)
Log(Cloud Computing Index)			100.5	104.3	54.65	-29.42
			(173.3)	(173.1)	(170.4)	(137.8)
Log(Data Storage Index)				-3,770	-3,350	-407.9
				(7,569)	(7,583)	(5,809)
Log(Digital Literacy Index)					34.31	33.42
					(191.6)	(157.2)
Log(Management Index)					281.1	292.2
					(370.3)	(328.0)
Log(Advertising Index)					198.4	135.7
					(123.7)	(97.20)
Lagged MV						0.259***
						(0.0656)
Constant	-5,015	-8,504	-8,734	-8,713	-9,772	-12,229
	(24,528)	(24,671)	(24,678)	(24,668)	(25,338)	(19,304)
Observations	23,935	23,935	23,935	23,935	23,935	23,896
Firm and Quarter FE	✓	✓	✓	✓	✓	✓

Table 7: SML Difference-in-Difference Regressions

Table Notes: *** p<0.01, ** p<0.05, * p<0.1. Standard errors clustered by firm in parentheses. The table reports results from estimating versions of equation 18 on a balanced panel of firms from the beginning of 2014 to the end of 2017. The logged indices, including human capital, are actually $log(x+1)$ in all cases. Columns are analogous to Table 5 for AI skills.

shocks in the same period) are the event study results following equation 19 with AI skills swapped for SML scores. Figure 5 shows no evidence of a contemporaneous market value shift during the time that TensorFlow is open-sourced, and toward the end of the sample period it appears that high SML firms begin to lose market value. That is, the firms with the most workers exposed to new ML capabilities are declining in their valuations. Firms with large proportions of their resources, including the workforce, exposed to new technologies are more likely to have to change their value creation processes in response to competitive pressure. Changing these processes can be very expensive, especially without in-house talent to implement new ones. These fixed costs of investment are simultaneously what insulate the rents of high value firms from competition ex-post and destroy the rents of incumbents at a discontinuity in technological capabilities. Just as fixed assets and intangible capital can provide upside exposure to employee skills, they can also create downside risks when adjustment is needed.

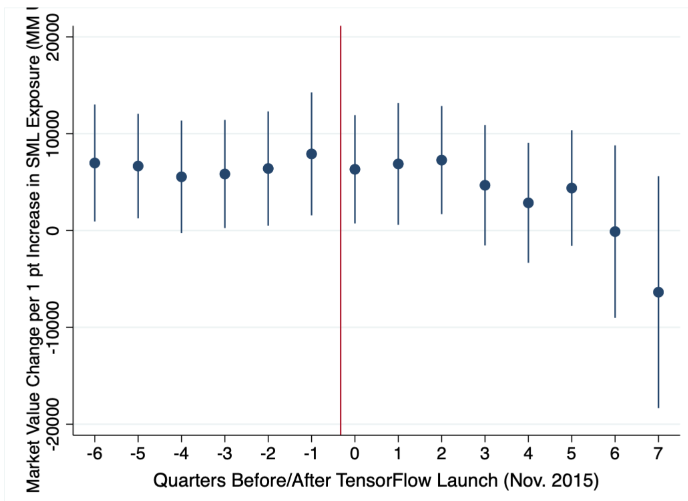

Figure 4: SML Event Study Results - Full Balanced Sample

Figure Notes: This plots the 95 percent confidence interval for the coefficients β_{4t} from equation 19 estimated on the set of covariates in column (5) of Table 7 over time. The upper and lower bounds of the 95 percent confidence interval are represented by vertical lines. The y-axis is the coefficient values for the regression of market value on the interacted time dummies and SML exposure. Standard errors are clustered by firm.

Having found little evidence for the total factor productivity or the workforce exposure hypothesis, I now turn to possible threats to identification of the price effect. The caveats on the main results are closely related to the worker productivity hypothesis as well, and will be addressed in section 7.

6.2 Threats to Identification and Causal Interpretations of Results

There are a number of possible threats to the validity of the assumptions necessary to make a causal inference from the AI skills difference-in-differences regressions. Though it seems to be the case that, in aggregate, there are no major pre-trends in market value per Figure 3, a more serious concern is selection issues across the sample. This is partly mitigated by using a balanced panel of firms

— by eliminating entry and exit into the sample there is a greater chance of recovering a causal estimate rather than a correlation reflecting compositional changes. But even within the balanced panel there may be changes in AI intensity that bias the coefficient estimates for the treatment effect. A related problem is that the market value behavior of firms without AI complements might be an inappropriate control for outcomes in companies that do have AI complements. It may also be the case that other skills, how the AI skills measure is coded (continuous versus discrete), or worker-level productivity effects (via the third hypothesis predicted by the theory) are driving the effects. In the subsequent section, I will test the robustness of the results to each of these factors.

There are, however, a few assumptions needed that are more difficult to test. One possible issue is that the interpretation of all of the foregoing analysis assumes that market participants are properly pricing the assets of publicly traded firms. That is, to some extent the market must be efficient with respect to pricing human capital-related intangible capital holdings. Eisfeldt and Papanikolaou (2013, 2014) provide a rationale and empirical investigation for why it seems to be the case that the market does price intangible assets, and human capital-related assets in particular. But for this specific case, how is it that market participants know how many AI workers are at each company? By 2020, numerous datasets with detailed worker data (similar to the ones used in this study) were readily available to hedge funds and other investors. These datasets existed, but were less prevalent, at the time of the TensorFlow shock. While verifying the information set of investors toward the end of 2015 would prove too difficult a challenge, a simpler constraint can be imposed that also helps deal with the compositional shift identification problem: breaking firms into quintiles (or any other quantile) reflecting their level of AI employment at the time of the TensorFlow launch. The sample is then consistent over time, and requires less from investors. Instead of knowing specific levels of AI skills by firm, they need only have a general sense of the level of AI use at each company. As it turns out, by the end of 2015 fewer than half of firms used AI. Figure **??** shows the histogram of firms by AI skills quintile. The second quintile is absorbed into the first with no skills, while the third quintile is small and also partially absorbed into the first. These quintiles will be applied for further analysis. I will also use a binary variable reflecting an indicator for hiring any AI skills or not to estimate a synthetic difference-in-differences model (Arkhangelsky et al., 2019) in section 7.

A final challenge for the identification of the TensorFlow effect is the assumption that it is TensorFlow causing these changes at all. Of course it is not possible to be absolutely certain that the TensorFlow launch caused the changes observed. The estimation is done at the quarterly level, which might include a variety of information shocks with respect to the AI content of firms. Abstraction is helpful here to some extent. Rather than considering the TensorFlow shock as unique to Google's decision, the real requirement to test whether or not worker potential AI skill acquisition is linked to the market value of employers is that there is a talent shock in that time period. Demonstrating the existence of a talent shock at that period in time is possible; it need not (only) come from TensorFlow. There are at least three separate datasets that corroborate the existence of a large talent shock at the time of the TensorFlow launch in late 2015: GitHub activity reflecting engagement with AI toolkits, Google Trends interest in deep learning and AI skillsets, and LinkedIn activity in adding

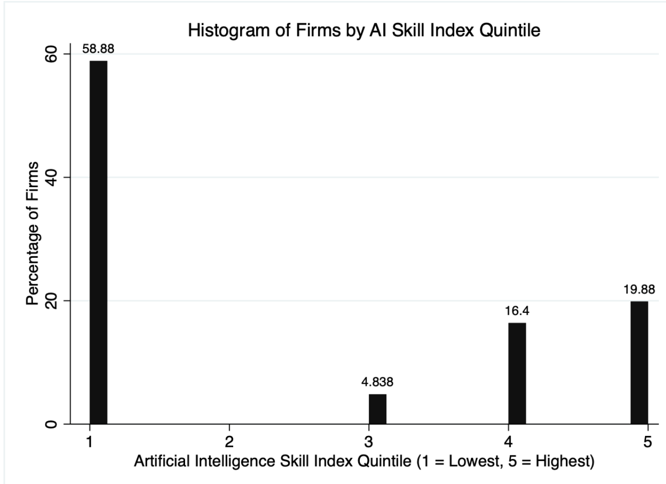

Figure 5: Histogram of Firms by AI Quintile

Figure Notes: This plot is a histogram of the percentage of firms in each AI Skill quintile. Most firms do not use AI, and the third quintile is small. The fourth and fifth quintiles are as expected. Differences from 20 percent membership there are due to ties in the discrete counts of skills.

deep learning and related skills. Figure 1(b) shows that deep learning skill additions increased at a rapid clip starting in late 2015. Using data from GitHub and Zhang et al. (2021), it is clear that engagement with AI software toolkits began to trend similarly around the same time. Figure 6 shows trends in GitHub stars, a measure of activity with respect to software modules.[32] The Google Trends data shows similar trends, with a rapid uptick in searches for TensorFlow (as might be expected) on the day of the launch. Figure 7 shows the relative search interest on Google Trends for deep learning, TensorFlow, and Theano (an older library) over time. Taken together, the three sources suggest that the talent shock did occur sometime around the launch of TensorFlow, even if it is possible that it were not TensorFlow causing the change.

7 Robustness Checks

7.1 Skills Placebos, Alternative AI Measures, and Synthetic Difference-in-Differences

How can the shifts in market value among AI using firms be better linked to AI skills and not other types of skills with related effects? Here we will test the hypothesis that worker productivity is influencing the market value of firms while also testing the price effect of AI. Skills placebos are an important check that can verify whether the level shifts in market value might be caused by other skills in the same time period, or perhaps if skills changes in general are covarying with firm value. Additionally these placebo checks provide some insight into the worker productivity hypothesis. AI workers have bundled skillsets in data science, cloud computing, and other techniques. If there was a shock to AI talent making it easier to acquire machine learning skills at the end of 2015, and any

[32]GitHub is a popular versioning control application and website. Software engineers can collaborate across open source projects, making contributions without conflicting versions. Additionally many companies, including Google, host open source projects on GitHub. GitHub was acquired by Microsoft in 2018.

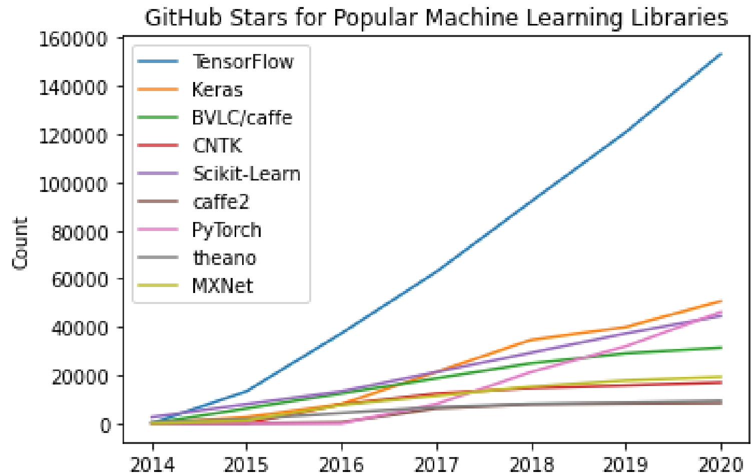

Figure 6: GitHub Stars over Time: ML Software

Figure Notes: This plot shows the GitHub star count for a number of machine learning libraries following Zhang et al. (2021) with data from GitHub. The data are publicly accessible via the 2021 AI Index Report: https://aiindex.stanford.edu/report/. TensorFlow is the most popular library by a wide margin according to this metric.

Figure 7: Google Trends Data for AI Skills

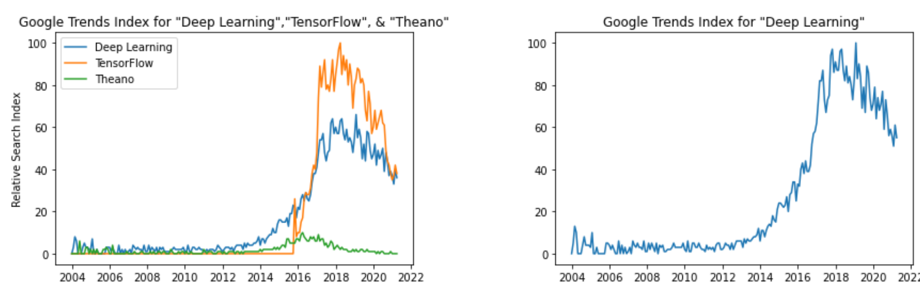

(a) AI Skills **(b) Deep Learning**

Figure Notes: This plot shows in panel (a) the relative search index for Deep Learning, TensorFlow, and Theano as search terms. Deep Learning accelerates in conjunction with TensorFlow after its launch, and prior to TensorFlow's launch it has little demonstrated interest. Theano interest briefly increases and declines in the same interval. Further development is not supported in Theano by its developing team. Panel (b) shows Deep Learning on its own.

response in market value were not due to AI alone, then we can expect a similar event study to show evidence of an increase in market value from those bundled skills as well (unless the worker appropriates the full value of those shifts).

Below I estimate event study coefficients of the form in equation 19. Instead of interacting AI skills with each of the time dummy variables, I instead interact high level skill cluster category indices in data science, management, and advertising with time dummies, and also low-level specific skill indices for linear regression, A/B Testing, and deep learning (specifically) to see if the observed market value effects from the previous sections are specific to AI-related talent. In each case, the goal is to be adversarial to the AI talent price hypothesis. Therefore the included controls for all specifications are lagged market value, total book value of assets, human capital years, and skills indices for cloud computing, digital literacy, data storage, big data skills, and AI. Additionally I include indices for data science, management, and advertising when those indices are not interacted with dummies. The results of these regressions are reported in Figure 8 below. To preserve an interpretation of a "per skill" value, these covariates are not logged. Therefore the effect under consideration should be a level shift at the time of the TensorFlow release. If something else shifts the market value of these skills in that time period, it will only occur once.

The placebos show that, for the most part, there is no shift in market value for the AI Skills acceleration start period (the final quarter of 2015) in response to skills other than deep learning. The broad categories show positive and statistically significant estimates for the value of Data Science, Management, and Advertising skills, but no major shifts immediately following the TensorFlow launch. A/B Testing and Linear Regression are two related competencies that are especially prevalent in the types of AI-intensive and data science-intensive employers that might see an increase in valuation because of TensorFlow. There appears to be a small increase in the point estimate for these skills following the TensorFlow launch, but it is not large enough to be out of the trend for these skill values. What this means, however, is that there is some (albeit weak) evidence for the worker productivity hypothesis because those related skill value point estimates do increase. It nevertheless seems unlikely to be the entire story. Deep Learning skills, in contrast, show a large and statistically significant increase immediately following the TensorFlow launch. At the same time, the confidence intervals for Deep Learning value in the pre-period are wider than they are for the post-period. This is because of massive additions of skills to the LinkedIn platform in this category. That increase in skill value due to deep learning, if not due to noise, would constitute a level shift in that period in the value of deep learning skills. It may indeed just be due to noise. This motivates the next section, where I analyze alternative codings of AI skills (quintiles and binary specifications).

Exposure to digital forms of capital is concentrated in some of the largest firms (Bessen, 2020; Tambe et al., 2020a). Given possible changes in the size of firms with respect to their AI talent flows, it makes sense to recast the AI measure into quintiles. These quintiles are set based on the AI skillset within each firm in the panel at the time of the TensorFlow launch. In some ways this is a more faithful measure of exposure for our outcome of choice; market participants are more likely to be able to discern AI using from non-AI using firms or quintiles of use than specific AI talent

Figure 8: Event Studies: Alternative Skill Placebos

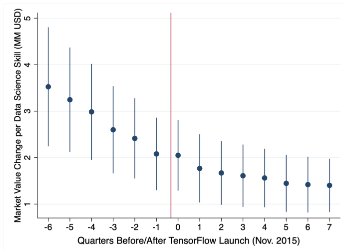

(a) Data Science Skills

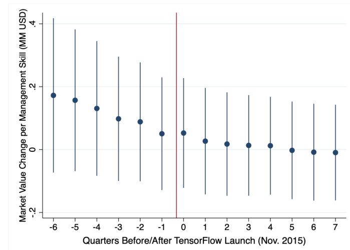

(b) Management Skills

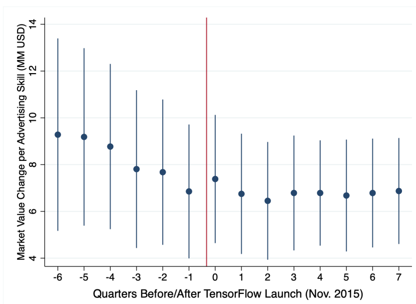

(c) Advertising Skills

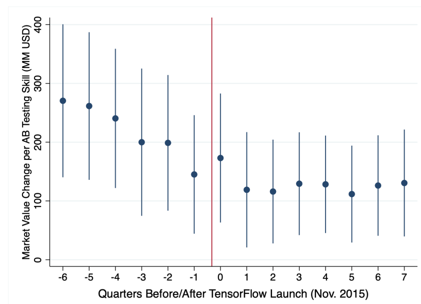

(d) A/B Testing Skills

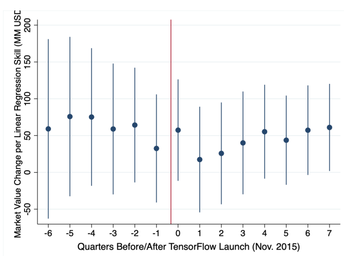

(e) Linear Regression Skills

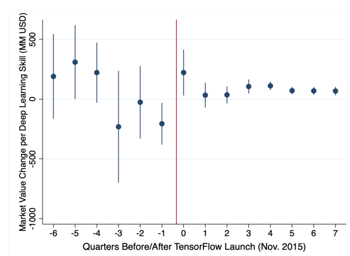

(f) Deep Learning Skills

Figure Notes: These charts show event study results per equation 19. Controls included are lagged market value, total book value of assets, years of education at the firm (human capital), and skills indices for cloud computing, digital literacy, data storage technologies, big data skills, AI skills, and advertising, management, and data science when those categories are not interacted with time fixed effects. Additionally firm and industry-time fixed effects are included. All skills are user-reported within publicly traded companies across the LinkedIn platform for the designated 2014-2017 (inclusive) time period. The TensorFlow launch in red is November 2015. Vertical bars are 95 percent confidence intervals. Standard errors clustered by Industry-Time where industry is defined as 4-Digit NAICS code.

on a continuous basis. Quintiles also reduce the possible effects of skewness in the skill measures and maintains a constant sample composition. The results for point estimates of these interacted dummies with quintiles for logged market value as an outcome variable with firm fixed effects are below in Figure 9. The entire table for logged and unlogged market value coefficients by quintile with firm and industry-time fixed effects are in the appendix in Table 12.

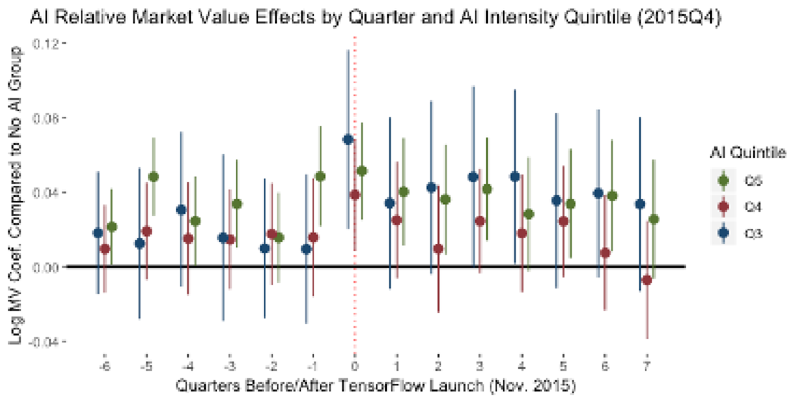

Figure 9: Quintile Log Market Value Event Study Results - Full Balanced Sample

Figure Notes: This plots the 95 percent confidence interval for the coefficients β_{4t} from equation 19 estimated on the set of covariates in column (5) of Table 5 over time (excluding AI skills). However, instead of using an AI skills index uses a fixed quintile of the AI Index as the AI treatment moderating variable. The upper and lower bounds of the 95 percent confidence interval are represented by vertical lines. The y-axis is the coefficient values for the regression of logged market value on the interacted time dummies and AI quintile exposure. The x-axis at zero constitutes the logged market value change of non-AI firms in quintiles 1 and 2. The green quintile is the most AI-intensive (5), then red (4), then blue (3). These coefficients therefore represent the change in percentage market value over time correlated with AI quintiles net of other skill indices. Standard errors are clustered by firm.

We see in Figure 9 a version of the effects in return space, as the market value outcome is logged. This reduces the influence of large market value firms in model estimation. The AI index has now been split to quintiles instead of its continuous version. Relative to the non-AI firm baseline in black at the 0 line, the only period in which all three quintile groups have a statistically significant and positive coefficient from the dummy variable interaction is in the time period immediately subsequent to the TensorFlow launch. Firms that already employ AI skillsets gain market value as a level shift upward relative to non-AI firms in this period. It is also apparent from this plot that, in green, the top AI-intensive firms are often outperforming their non-AI counterparts. That suggests looking into 1) coefficient estimates on the original specification without Top Quintile firms, and 2) that alternative control group specifications may be informative as to the source of effects (and if large firms are especially exposed to AI value boons). Table 8 below re-estimates the results in Table 5, but this time excluding the top quintile of firms.

The coefficient estimates are still broadly indicative of the same trends without the top quintile, but with considerably smaller coefficients. Now the putative effects of the TensorFlow launch are between \$302 million (standard error \$158 million) in column (6) and \$333 million (standard error \$173 million) in column (3) per doubling of AI talent. These are only statistically significant at a 10 percent size. This suggests that a large component of the valuation changes observed in earlier

AI Difference-in-Difference, No Top Quintile	(1) Market Value	(2) Market Value	(3) Market Value	(4) Market Value	(5) Market Value	(6) Market Value
Lagged MV						0.0799*
						(0.0464)
Total Assets	1.015***	1.015***	1.015***	1.015***	1.015***	0.968***
	(0.0117)	(0.0118)	(0.0118)	(0.0118)	(0.0118)	(0.0261)
Log(Human Capital)	578.8***	578.4***	583.0***	586.4***	591.4***	513.5***
	(104.4)	(105.1)	(106.9)	(106.7)	(109.2)	(108.3)
Log(AI Index)	-331.4	-332.4	-326.9	-321.6	-318.8	-319.0
	(212.3)	(213.6)	(215.3)	(214.9)	(214.3)	(201.6)
Log(AI Index)xPostTF	326.5*	325.9*	333.4*	330.8*	325.4*	302.0*
	(173.3)	(173.7)	(172.8)	(172.5)	(172.4)	(158.1)
Log(Data Science Index)		11.49	29.75	36.41	13.28	-0.300
		(112.6)	(105.1)	(104.3)	(104.4)	(97.60)
Log(Cloud Computing Index)			-82.79	-74.31	-94.98	-102.4
			(124.0)	(123.7)	(125.8)	(117.1)
Log(Data Storage Index)				-8,027*	-7,489*	-6,163*
				(4,096)	(3,967)	(3,607)
Log(Digital Literacy Index)					-89.04	-101.5
					(112.4)	(103.7)
Log(Management Index)					64.37	117.4
					(166.1)	(151.0)
Log(Advertising Index)					171.0*	142.1
					(97.48)	(89.48)
Constant	-2,938***	-2,969***	-2,948***	-2,878***	-3,150**	-3,030***
	(1,042)	(1,019)	(1,015)	(1,013)	(1,228)	(1,150)
Observations	16,961	16,961	16,961	16,961	16,961	16,959
Firm and Quarter FE	✓	✓	✓	✓	✓	✓

Table 8: AI Skills Difference-in-Difference Regressions without Top Quintile

Table Notes: *** p<0.01, ** p<0.05, * p<0.1. Standard errors clustered by firm in parentheses. The table reports results from estimating versions of equation 18 on a balanced panel of firms from the beginning of 2014 to the end of 2017. The logged indices, including human capital, are actually $log(x+1)$ in all cases. Columns are identical to Table 5 for AI skills, but exclude the top quintile of AI firms.

sections of the paper was driven by large firms. Since IT assets generally are concentrated in some of the highest market value firms and the link between IT, productivity, and rents is well-established (Brynjolfsson and Hitt, 2000; Corrado et al., 2009; Tambe and Hitt, 2012; Tambe, 2014; McGrattan, 2017; Tambe et al., 2020a), this is reasonably predicted by the literature as a likelihood for AI as well. Figure 10 plots the equivalent event study design as in Figure 3, this time dropping the top quintile as well. The effects are noticeably smaller for non-top firms.

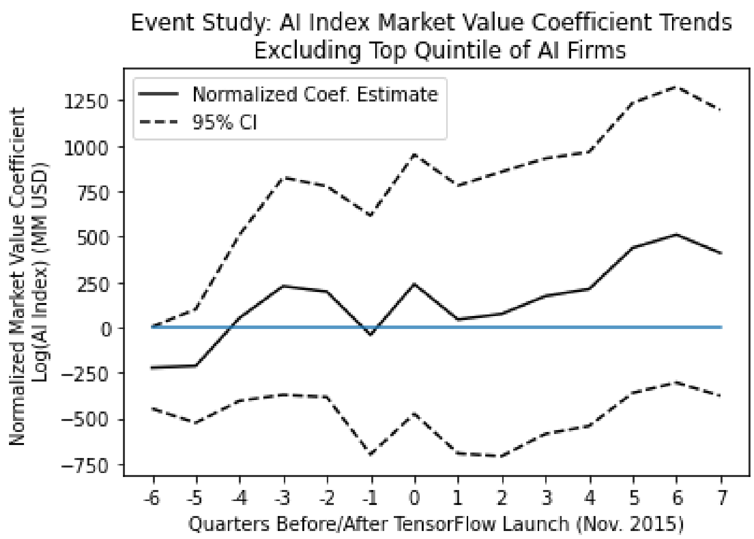

Figure 10: Event Study Results - No Top Quintile Balanced Sample

Figure Notes: This plots the 95 percent confidence interval for the coefficients β_{4t} from equation 19 estimated on the set of covariates in column (5) of Table 8 over time. All values have been normalized by subtracting the pre-TensorFlow mean value of the coefficients to set an (average) baseline value of 0 in the pre-treatment period. The zero benchmark represents firms that do not use AI at all. The upper and lower bounds of the 95 percent confidence interval are represented with dashed lines, and the 0 line is in blue. The y-axis is the coefficient values for the regression of market value on the interacted time dummies and AI skills. A coefficient of 1000, for example, means that a 100 percent increase in AI in that period predicts a $1 billion increase in market value. Standard errors are clustered by firm. Omits the top quintile of AI index firms.

The diminished estimates from excluding the top quintile are indicative that the results may be sensitive to choice of control group. In comparing firms with AI talent (and therefore AI-related complementary assets) to firms without AI talent, the aim is to estimate the effect of an AI talent shock between groups that are exposed and those that are not. The ideal experiment, however, is to think about what the same firm might do in response to the talent change without any AI talent versus what occurs in reality. For AI-intensive firms, we would like to compare to a doppelganger firm that does not have any AI skills. We can approximate this thought experiment in observational data with synthetic difference-in-differences (Arkhangelsky et al., 2019), a modification on the standard difference-in-difference framework that combines two-way fixed effects with the insights from the synthetic control literature. In synthetic control methods, pre-treatment trends are aligned by finding weights over untreated units that synthetically recreate the outcomes of treated units (and omits unit fixed effects). Synthetic controls, however, are not amenable as a method to large panels. In synthetic difference-in-differences, there are unit weights on the standard difference-in-differences two-way fixed effect squared errors as in synthetic controls, but there are also time weights. Since the difference in pre-treatment control and treatment group averages changes over time in the pre-period and this difference over time might be predictive of differences in the post-period, synthetic difference-in-differences reweights by time as well. This has the benefit of mitigating the biasing effect of time periods in the pre-period that are substantially different than the post-period that might occur with synthetic controls. Synthetic difference-in-differences also includes a unit fixed effect, which in the case of firm-level analyses constitutes a large component of the overall variation. Per Arkhangelsky et al. (2019), the synthetic difference-in-difference solves:

$$(\hat{\beta}_1^{sdid}, \hat{\beta}_0, \hat{\mu}_i, \hat{\nu}_t) = \underset{\beta_0,\mu_i,\nu_t,\beta_1}{\arg\min}\, [\sum_{i=1}^{N}\sum_{t=1}^{T}(Y_{it} - \beta_0 - \mu_i - \nu_t - \beta_1 X_{it})^2 \hat{\omega}_i^{sdid} \hat{\lambda}_t^{sdid}] \tag{21}$$

In this case, X_{it} is a binary treatment variable that I use to represent having nonzero AI talent or not (encoded as 1 if AI skills are present, else 0), μ_i and ν_t are unit and time fixed effects, respectively, β_0 is an intercept, and β_1 represents the treatment effect of interested (a shock to AI Talent). $\hat{\omega}_i^{sdid}$ and $\hat{\lambda}_t^{sdid}$ are the unit weights (as in synthetic controls) and the time weights that enable a more robust treatment effect. In other words, this method will better compare similar units with and without AI in similar time periods. Standard errors are calculated with a jackknife estimation procedure described in Arkhangelsky et al. (2019) and implemented in the associated R package. Effectively this process allows for better modeling of pre-treatment outcomes. Results for AI and other selected binary covariates are below in Table 9.

Skill Group	Estimate	Standard Error	N Treated
Artificial Intelligence	1.6	0.75	605
Data Science	0.72	0.41	1262
Deep Learning	14.3	4.9	60
Cloud Computing	0.93	0.53	943
Business Management	0.25	0.39	1357
Linear Regression	3.06	1.74	199
SML (Median)	0.16	0.3	1037

Table 9: Synthetic Difference-in-Difference Results

Table Notes: Jackknife standard errors reported in the table. The table reports results from estimating versions of equation 21 on a balanced panel of firms from the beginning of 2014 to the end of 2017. All skill indices (AI, data science, cloud computing, business management, and linear regression) are recast as 1 if skill values are nonzero at the time of the TensorFlow launch and 0 otherwise. SML is made into a binary variable by recoding it as 1 if equal to or above median and 0 if below. Estimates are in billions USD.

In Table 9, the only statistically significant estimates at the 5 percent level are those for AI skills and Deep Learning skills. With better matching of treatment and control (via the more flexible synthetic DID framework), the coefficient estimates actually get larger. Given the identification assumptions hold, Being a firm with some AI workers at the time of the TensorFlow launch now causes an approximate \$1.6 billion (standard error \$750 million) increase in market value. If the firm is using deep learning in particular, the estimated caused increase in market value is $14.3 billion (standard error $4.9 billion). Notably there not many firms using deep learning (only 60), so this value increase only affects a few firms who are lucky enough to be positioned for it. AI is more widespread at the time, with 605 firms in the AI-intensive group. The other point estimates are all positive, but not statistically significant. The same placebo checks as before therefore hold, and the composition of the sample is held constant in this analysis. With this better matching of an AI-intensive firm to a synthetic non-AI using comparison group, we see treatment effects are larger. Taken together, it appears that the AI talent shock (very likely related to TensorFlow at a minimum) caused a revaluation of firms with AI-complementary assets per the pricing hypothesis. The two flavors of productivity hypotheses did not show up in valuation for other skills to a significant extent and the SML exposure results suggest negative value effects, if any at all.

8 Discussion and Managerial Implications

Through a number of results and regression specifications, the balance of the evidence suggests that the open source launch of TensorFlow (or a concurrent shock making it easier to acquire AI-related skills) caused an increase in the market value of firms with AI complements. This change likely occurred via the hypothesized price effect: sunk fixed investments at firms with firm-specific value became more valuable as the costs to service complementary investments in talent dropped. The advantage of using market value is that it is forward-looking, as opposed to many other important measures like productivity which can only be measured contemporaneously. There is little observable effect of the TensorFlow launch on total factory productivity (proxied for via a production function OLS estimation and by value of other skills), and the productivity effects for talented workers with complementary skills to AI appear muted at best. Lastly, firms with lots of exposure to applications of AI and machine learning, as measured by the weighted SML measure, appear to be either unaffected or possibly negatively affected with a lag by the shift in skill acquisition costs. Returning to the price effect is therefore a diagnosis of exclusion in a way — the other possibilities are ruled out.

These results are robust to a variety of placebo checks and other diagnostics. The parallel trends assumption can of course be problematic for any difference-in-differences model. A visualization of pre-trends yields no obvious problem, but firms with AI complements do tend to outperform non-AI counterparts in the pre-period. Dropping these early adopters leads to diminished estimates that are qualitatively similar. But additionally there is the question of whether or not non-AI firms as a group constitute an appropriate comparison group. Using the synthetic difference-in-difference method to build a better comparison group, the effects of the rapid increase in ML talent starting at the end of 2015 is even stronger for companies that have started up the AI and deep learning adoption curve. This type of model, in recasting the AI exposure variable as binary, further handles a separate problem: how do investors know how intensively companies use AI? It is far easier to tell the extensive margin of technology use as an onlooker than the intensive margin. The results are also robust to sample selection changes that might be introduced by using a continuous measure, even if the panel is balanced. When looking at quintiles that are fixed, the results are qualitatively similar to the continuous measures. These quintiles reveal as well that the largest effects are concentrated. Only 60 firms are using deep learning at the time of the TensorFlow launch, and the point estimates for their market value growth as a result are large. Firms in the third and fourth quintile of AI talent also appear to receive a market value boost from more abundant expected future talent, but that boost is short-lived in comparison to the larger AI skill index firms. Like other forms of IT, the market value gains to AI investment seem to accrue to a relatively concentrated set of firms. The rapid proliferation of skillsets, however, is spread out over many different workers as interest in deep learning skills has climbed rapidly in recent years.

There are many facets to the managerial implications of these results. The first, as indicated by the theory, is that firm-specific capital investments provide exposure to future employee skillsets. Employee skills have different risk-reward profiles. Technology skill value can be highly volatile, with quick adoption and abandonment (Deming and Noray, 2018; Horton and Tambe, 2019). The firm-specific technological assets in a company will grow in value if the human capital complements become more abundant. Even firms, as Hal Varian suggested, can become expensive complements to something that is rapidly becoming cheaper. Those technological assets can easily become obsolete or impaired by new technologies. Companies with highly ML-exposed employees are declining in market value over the last few quarters of the skill sample relative to peers. Managers can therefore consider employee skills as part of their investment portfolio, and even guide them deliberately as companies like Google, Facebook, Uber, IBM, and LinkedIn have by making open source projects a key part of their business strategy. The choice to use OSS may avoid future training costs. In the past, as with the development of MapReduce, Google did not open source its software. Later when workers were trained to use Hadoop for similar purposes, Google would have to invest in retraining their new hires to use internal systems. Closed systems therefore carry their own costs. Additionally from a measurement perspective, keeping a careful inventory of the human capital varieties within firms will be occasionally be important to understanding what is driving the share price. Changes like the deep learning talent effects are infrequent, but have the potential to be large.[33] New large-scale labor platforms like LinkedIn and GitHub can help firms take inventory of their skillsets. Finally, though this paper considers technological skills in detail, such changes are possible as well outside of technology skills. New software packages are relatively easy to measure in comparison to social skills or leadership ability, but they are not uniformly more valuable. Investing in measurement of non-technological human capital may be a high return activity.

[33] When Andrew Ng left Baidu, for example, the market capitalization of Baidu dropped by about \$1.5 billion: https://www.businessinsider.com/baidu-value-took-a-15-billion-plunge-after-chief-scientist-andrew-ng-announced-hes-leaving-2017-3

9 Conclusion

This paper investigates the effect of making technological talent in Artificial Intelligence more abundant via a shock to the worker costs of investing in machine learning skills. The results indicate that AI talent-exposed firms, having sunk firm-specific investments in AI, experienced larger growth as a result of Google's decision to open source TensorFlow (or a similar shock causing a growth in AI talent at the same time). The results are robust to a variety of specifications, placebo checks, and additional analyses.

This study is among the first to analyze a large scale online database of the supply-side of AI talent in firms to understand the effects of AI adoption in companies. I also link human capital acquisition decisions made by employees to the market values of their employers, finding that the price effects at the outset of a skill proliferation event are the most likely explanation and not contemporaneous productivity increases or prior overall workforce exposure as an intangible asset. The market value increases are more likely due to the value of AI projects increasing instead. Technology skills like deep learning are particularly well-suited to analyses of this kind because coverage of tech workers on online platforms tends to be high, their skills are easily captured in online databases and explicitly named (and sometimes quite strangely named), and the skills can arise and die off quickly. Nevertheless, measurement can be subject to a host of biases and issues when using online platforms. While the measurement is challenging, the economics and managerial implications are relatively simple: firm-specific assets provide exposure to the employee human capital complements. The q-value of companies is linked to the technological knowledge they rent from workers. When workers decide to pursue new skills, it is a win-win for employers as well.

References

Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., Devin, M., Ghemawat, S., Irving, G., Isard, M., et al. (2016). Tensorflow: A system for large-scale machine learning. In *OSDI*, volume 16, pages 265–283.

Abis, S. and Veldkamp, L. (2020). The changing economics of knowledge production. *Available at SSRN 3570130*.

Acemoglu, D. and Autor, D. (2011). Skills, tasks and technologies: Implications for employment and earnings. In *Handbook of Labor Economics*, volume 4, pages 1043–1171.

Acemoglu, D., Autor, D., Hazell, J., and Restrepo, P. (2020). AI and jobs: Evidence from online vacancies. Technical report, National Bureau of Economic Research.

Acemoglu, D. and Pischke, J.-S. (1998). Why do firms train? Theory and evidence. *The Quarterly Journal of Economics*, 113(1):79–119.

Acemoglu, D. and Pischke, J.-S. (1999). Beyond becker: Training in imperfect labour markets. *The Economic Journal*, 109(453):112–142.

Acemoglu, D. and Restrepo, P. (2018). The race between machine and man: Implications of technology for growth, factor shares and employment. *American Economic Review*.

Agrawal, A., Gans, J. S., and Goldfarb, A. (2017). Prediction, judgment, and uncertainty. In *Economics of Artificial Intelligence*. University of Chicago Press.

Agrawal, A., Gans, J. S., and Goldfarb, A. (2018). *Prediction Machines: The Simple Economics of Artificial Intelligence*. Harvard Business Review Press, Boston.

Agrawal, A., Goldfarb, A., and Teodoridis, F. (2016). Understanding the changing structure of scientific inquiry. *American Economic Journal: Applied Economics*, 8(1):100–128.

Alekseeva, L., Azar, J., Gine, M., Samila, S., and Taska, B. (2020). The demand for AI skills in the labor market.

Andrews, D., Criscuolo, C., Gal, P., et al. (2015). *Frontier Firms, Technology Diffusion and Public Policy: Micro Evidence from OECD Countries*, volume 2. OECD Publishing.

Arkhangelsky, D., Athey, S., Hirshberg, D. A., Imbens, G. W., and Wager, S. (2019). Synthetic difference in differences. Technical report, National Bureau of Economic Research.

Autor, D., Levy, F., and Murnane, R. J. (2003). The skill content of recent technological change: An empirical exploration. *The Quarterly Journal of Economics*, 118(4):1279–1333.

Autor, D. H. (2014). Polanyi's paradox and the shape of employment growth. *Working Paper*, pages 129–178.

Autor, D. H. and Dorn, D. (2013). The growth of low-skill service jobs and the polarization of the US Labor Market. *American Economic Review*, 103(5):1553–1597.

Azar, J. A., Marinescu, I., Steinbaum, M. I., and Taska, B. (2018). Concentration in US labor markets: Evidence from online vacancy data. Technical report, National Bureau of Economic Research.

Babina, T., Fedyk, A., He, A. X., and Hodson, J. (2020). Artificial intelligence, firm growth, and industry concentration. *Firm Growth, and Industry Concentration (November 22, 2020)*.

Balasubramanian, N., Chang, J. W., Sakakibara, M., Sivadasan, J., and Starr, E. (2018). Locked in? The enforceability of covenants not to compete and the careers of high-tech workers.

Bardhan, I., Krishnan, V., and Lin, S. (2013). Research note—business value of information technology: Testing the interaction effect of it and r&d on tobin's q. *Information Systems Research*, 24(4):1147–1161.

Barua, A., Konana, P., Whinston, A. B., and Yin, F. (2004). An empirical investigation of net-enabled business value. *MIS quarterly*, pages 585–620.

Barua, A., Sophie Lee, C., and Whinston, A. B. (1996). The calculus of reengineering. *Information Systems Research*, 7(4):409–428.

Becker, G. S. (1962). Investment in human capital: A theoretical analysis. *Journal of Political Economy*, 70(5, Part 2):9–49.

Bedard, K. (2001). Human capital versus signaling models: university access and high school dropouts. *Journal of Political Economy*, 109(4):749–775.

Benzell, S., Lagarda, G., and Rock, D. (2018a). Do labor demand shifts occur within firms or across them? non-routine-biased technological change, 2000-2016. *Unpublished Working Paper*.

Benzell, S., Lagarda, G., and Rock, D. (2018b). Skill-biased technological change within and across firms: 2000-2016.

Bessen, J. (2020). Industry concentration and information technology. *The Journal of Law and Economics*, 63(3):531–555.

Bessen, J. E., Impink, S. M., Reichensperger, L., and Seamans, R. (2020). GDPR and the importance of data to AI startups. *Available at SSRN 3576714*.

Bharadwaj, A. S. (2000). A resource-based perspective on information technology capability and firm performance: An empirical investigation. *MIS quarterly*, pages 169–196.

Bharadwaj, A. S., Bharadwaj, S. G., and Konsynski, B. R. (1999). Information technology effects on firm performance as measured by Tobin's q. *Management science*, 45(7):1008–1024.

Bresnahan, T. (2010). General purpose technologies. *Handbook of the Economics of Innovation*, 2(1):761–791.

Bresnahan, T., Greenstein, S., Brownstone, D., and Flamm, K. (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., Brynjolfsson, E., and Hitt, L. M. (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 Trajtenberg, M. (1995). General purpose technologies 'Engines of growth'? *Journal of Econometrics*, 65(1):83–108.

Brynjolfsson, E., Frank, M. R., Mitchell, T., Rahwan, I., and Rock, D. (2019). Quantifying the Distribution of Machine Learning's Impact on Work.

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

Brynjolfsson, E., Hitt, L. M., and Kim, H. H. (2011). Strength in numbers: How does data-driven decisionmaking affect firm performance? *Available at SSRN 1819486*.

Brynjolfsson, E., Hitt, L. M. L. M., and Yang, S. (2002). Intangible assets: Computers and organizational capital. *Brookings Papers on Economic Activity*, 2002(1):137–198.

Brynjolfsson, E., Hui, X., and Liu, M. (2018a). Does machine translation affect international trade? Evidence from a large digital platform. Technical report, National Bureau of Economic Research.

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

Brynjolfsson, E. and Mitchell, T. (2017). What can machine learning do? Workforce implications. *Science*, 358(6370):1530–1534.

Brynjolfsson, E., Mitchell, T., and Rock, D. (2018b). What can machines learn, and what does it mean for occupations and the economy? *AEA Papers and Proceedings*, pages 43–47.

Brynjolfsson, E., Rock, D., and Syverson, C. (2018c). Artificial intelligence and the modern productivity paradox: A clash of expectations and statistics. In *Economics of Artificial Intelligence*. University of Chicago Press.

Brynjolfsson, E., Rock, D., and Syverson, C. (2021). The productivity J-curve: How intangibles complement general purpose technologies. *American Economic Journal: Macroeconomics*, 13(1):333–72.

Bughin, J., Hazan, E., Ramaswamy, S., Chui, M., Allas, T., Dahlstrom, P., Henke, N., and Trench, M. (2017). Artificial intelligence – the next digital frontier? *Artificial Intelligence*, page 80.

Caldwell, S. and Danieli, O. (2018). Outside options in the labor market. *Unpublished manuscript*.

Chollet, F. (2015). Keras: Deep learning library for theano and tensorflow. *URL: https://keras.io/k*, 7(8).

Choudhury, P., Starr, E., and Agarwal, R. (2020). Machine learning and human capital complementarities: Experimental evidence on bias mitigation. *Strategic Management Journal*, 41(8):1381–1411.

Claussen, J., Peukert, C., and Sen, A. (2019). The editor vs. the algorithm: Targeting, data and externalities in online news. *Data and Externalities in Online News (June 5, 2019)*.

Cockburn, I. M., Henderson, R., and Stern, S. (2018). The impact of artificial intelligence on innovation: An exploratory analysis. *National Bureau of Economic Research*, No. w24449(September).

Corrado, C., Hulten, C., and Sichel, D. (2009). Intangible capital and U.S. economic growth. 55(3):661–685.

Cowgill, B. (2018). Bias and productivity in humans and algorithms: Theory and evidence from resume screening. *Columbia Business School, Columbia University*, 29.

Crouzet, N. and Eberly, J. (2018). Intangibles, investment, and efficiency. *AEA Papers and Proceedings*.

Crouzet, N. and Eberly, J. C. (2019). Understanding weak capital investment: The role of market concentration and intangibles. Technical report, National Bureau of Economic Research.

Cummins, J. G. (2005). A new approach to the valuation of intangible capital. In *Measuring Capital in the New Economy*, pages 47–72. University of Chicago Press.

Deming, D. J. and Noray, K. L. (2018). STEM careers and technological change.

Dewan, S. and Kraemer, K. L. (2000). Information technology and productivity: Evidence from country-level data. *Management science*, 46(4):548–562.

Dewan, S., Michael, S. C., and Min, C.-k. (1998). Firm characteristics and investments in information technology: Scale and scope effects. *Information Systems Research*, 9(3):219–232.

Dewan, S., Shi, C., and Gurbaxani, V. (2007). Investigating the risk–return relationship of information technology investment: Firm-level empirical analysis. *Management science*, 53(12):1829–1842.

Eisfeldt, A. L. and Papanikolaou, D. (2013). Organization capital and the cross-section of expected returns. *Journal of Finance*, 68(4):1365–1406.

Eisfeldt, A. L. and Papanikolaou, D. (2014). The value and ownership of intangible capital. *American Economic Review*, 104(5):189–94.

Erel, I., Stern, L. H., Tan, C., and Weisbach, M. S. (2018). Selecting directors using machine learning. Technical report, National Bureau of Economic Research.

Ewens, M., Nanda, R., and Rhodes-Kropf, M. (2018). Cost of experimentation and the evolution of venture capital. *Journal of Financial Economics*, 128(3):422–442.

Farboodi, M., Mihet, R., Philippon, T., and Veldkamp, L. (2019). Big data and firm dynamics. In *AEA Papers and Proceedings*, volume 109, pages 38–42.

Felten, E. W., Raj, M., and Seamans, R. (2018). A method to link advances in artificial intelligence to occupational abilities. *AEA Papers and Proceedings*, 108:54–57.

Fitoussi, D. and Gurbaxani, V. (2012). IT outsourcing contracts and performance measurement. *Information Systems Research*, 23(1):129–143.

Frank, M. R., Autor, D., Bessen, J. E., Brynjolfsson, E., Cebrian, M., Deming, D. J., Feldman, M., Groh, M., Lobo, J., Moro, E., et al. (2019). Toward understanding the impact of artificial intelligence on labor. *Proceedings of the National Academy of Sciences*, 116(14):6531–6539.

Furman, J. and Seamans, R. (2018). AI and the economy. pages 1–33.

Fuster, A., Goldsmith-Pinkham, P., Ramadorai, T., and Walther, A. (2020). Predictably unequal? the effects of machine learning on credit markets. *The Effects of Machine Learning on Credit Markets (October 1, 2020)*.

Goldfarb, A., Taska, B., and Teodoridis, F. (2020). Could machine learning be a general purpose technology? A comparison of emerging technologies using data from online job postings.

Greenstein, S. and Nagle, F. (2014). Digital dark matter and the economic contribution of Apache. *Research Policy*, 43(4):623–631.

Grennan, J. and Michaely, R. (2019). Artificial intelligence and the future of work: Evidence from analysts.

Hall, B. H. (1993). The Stock market's valuation of R&D investment during the 1980's. *The American Economic Review*, 83(2):259–264.

Hall, B. H. (2006). R&D, productivity, and market value. *IFS Working Papers, Institute for Fiscal Studies (IFS) 06/23*.

Hall, R. E. (2001). The stock market and capital accumulation. *The American Economic Review*, 91(5):1185–1202.

Hall, R. E. (2004). Measuring factor adjustment costs. *The Quarterly Journal of Economics*, 119(3):899–927.

Hall, R. E. (2017). High discounts and high unemployment. *American Economic Review*, 107(2):305–30.

Haskel, J. and Westlake, S. (2017). *Capitalism without Capital: The Rise of the Intangible Economy*. Princeton University Press.

Hayashi, F. (1982). Tobin's marginal q and average q: A neoclassical interpretation. *Econometrica*, 50(1):213–224.

Hippel, E. and Krogh, G. (2003). Open source software and the "private-collective" innovation model: Issues for organization science. *Organization science*, 14(2):209–223.

Horton, J. J. and Tambe, P. (2019). The death of a technical skill. *Unpublished Manuscript*.

Iansiti, M. and Lakhani, K. R. (2020). *Competing in the Age of AI: Strategy and Leadership When Algorithms and Networks Run the World*. Harvard Business Press.

Jeffers, J. S. (2017). The impact of restricting labor mobility on corporate investment and entrepreneurship.

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

Jones, C. I. and Tonetti, C. (2020). Nonrivalry and the economics of data. *American Economic Review*, 110(9):2819–58.

Kaldor, N. (1966). Marginal productivity and the macro-economic theories of distribution: Comment on Samuelson and Modigliani. *The Review of Economic Studies*, 33(4):309–319.

Kokkodis, M. (2019). Skillset diversification in online labor markets: Reputation losses and opportunity gains. *Available at SSRN 3444363*.

Kokkodis, M. and Ipeirotis, P. G. (2016). Reputation transferability in online labor markets. *Management Science*, 62(6):1687–1706.

Lakani, K. and Hippel, E. (2002). How open source software works:'free'user-to user assistance'. *Research Policy*, 32(2003):923–943.

LeCun, Y., Bengio, Y., and Hinton, G. (2015). Deep learning. *Nature*, 521(7553):436–444.

LeCun, Y., Bottou, L., Bengio, Y., and Haffner, P. (1998). Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2323.

Li, D., Raymond, L. R., and Bergman, P. (2020). Hiring as exploration. Technical report, National Bureau of Economic Research.

Lucas, R. E. (1967). Adjustment costs and the theory of supply. *Journal of Political Economy*, 75(4, Part 1):321–334.

Lustig, H., Syverson, C., and Van Nieuwerburgh, S. (2011). Technological change and the growing inequality in managerial compensation. *Journal of Financial Economics*, 99(3):601–627.

Mani, D., Barua, A., and Whinston, A. (2010). An empirical analysis of the impact of information capabilities design on business process outsourcing performance. *Mis Quarterly*, pages 39–62.

Mani, D., Barua, A., and Whinston, A. B. (2012). An empirical analysis of the contractual and information structures of business process outsourcing relationships. *Information Systems Research*, 23(3-part-1):618–634.

Marrano, M. G., Haskel, J., and Wallis, G. (2009). What happened to the knowledge economy? ICT, intangible investment, and Britain's productivity record revisited. 55(3):686–716.

Marx, M. (2011). The firm strikes back: Non-compete agreements and the mobility of technical professionals. *American Sociological Review*, 76(5):695–712.

McGrattan, E. R. (2017). Intangible capital and measured productivity. Technical report, National Bureau of Economic Research.

Melville, N., Kraemer, K., and Gurbaxani, V. (2004). Information technology and organizational performance: An integrative model of it business value. *MIS Quarterly*, 28(2):283–322.

Miric, M. and Ozalp, H. (2020). Standardized tools and the generalizability of human capital: The impact of standardized technologies on employee mobility. *Available at SSRN 3554224*.

Mithas, S., Ramasubbu, N., and Sambamurthy, V. (2011). How information management capability influences firm performance. *MIS quarterly*, pages 237–256.

Mithas, S., Tafti, A., Bardhan, I., and Goh, J. M. (2012). Information technology and firm profitability : Mechanisms and empirical evidence. *MIS Quarterly*, 36(1):205–224.

Mueller, H., Groger, A., Hersh, J., Matranga, A., and Serrat, J. (2020). Monitoring war destruction from space: A machine learning approach. *arXiv preprint arXiv:2010.05970*.

Nagaraj, A. (2021). Information seeding and knowledge production in online communities: Evidence from OpenStreetMap. *Management Science*.

Nagle, F. (2019). Open source software and firm productivity. *Management Science*, 65(3):1191–1215.

Nicolaus Henke, Jacques Bughin, Michael Chui, James Manyika, Tamim Saleh, Bill Wiseman, and Guru Sethupathy (2016). The age of analytics: Competing in a data-driven world. Technical report, Technical report, San Francisco: McKinsey & Company.

Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., Lin, Z., Desmaison, A., Antiga, L., and Lerer, A. (2017). Automatic differentiation in pytorch.

Peters, R. H. and Taylor, L. A. (2017). Intangible capital and the investment-q relation. *Journal of Financial Economics*, 123(2):251–272.

Polanyi, M. (1966). The logic of tacit inference. *Philosophy*, 41(155):1–18.

Raj, M. and Seamans, R. (2018). AI, labor, productivity, and the need for firm-level data. In *Economics of Artificial Intelligence*. University of Chicago Press.

Roach, M. and Sauermann, H. (2010). A taste for science? PhD scientists' academic orientation and self-selection into research careers in industry. *Research Policy*.

Rosen, S. (1981). The economics of superstars. *The American economic review*, 71(5):845–858.

Rumelhart, D. E., Hinton, G. E., and Williams, R. J. (1986). Learning representations by back-propagating errors. *Nature*, 323(6088):533–536.

Sambamurthy, V., Bharadwaj, A., and Grover, V. (2003). Shaping agility through digital options: Reconceptualizing the role of information technology in contemporary firms. *MIS quarterly*, pages 237–263.

Saunders, A. and Brynjolfsson, E. (2016). Valuing information technology related intangible assets. *Mis Quarterly*, 40(1).

Saunders, A. and Tambe, P. (2015). Data assets and industry competition: Evidence from 10-K filings. (September).

Schubert, G., Stansbury, A., and Taska, B. (2020). Employer concentration and outside options. Technical report, mimeo, Harvard University.

Spence, M. (1978). Job market signaling. In *Uncertainty in Economics*, pages 281–306. Elsevier.

Starr, E. (2015). Consider this: Firm-sponsored training and the enforceability of covenants not to compete. *Working Paper*.

Starr, E., Balasubramanian, N., and Sakakibara, M. (2017). Screening spinouts? How noncompete enforceability affects the creation, growth, and survival of new firms. *Management Science*, 64(2):552–572.

Stern, S. (2004). Do scientists pay to be scientists? *Management science*, 50(6):835–853.

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

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

Tambe, P. and Hitt, L. M. (2012). The productivity of information technology investments: New evidence from IT labor data. *Information Systems Research*.

Tambe, P., Ye, X., and Cappelli, P. (2020b). Paying to program? Engineering brand and high-tech wages. *Management Science*, 66(7):3010–3028.

Teodoridis, F. (2017). Understanding team knowledge production: The interrelated roles of technology and expertise. *Management Science*.

Thompson, N. (2017). The economic impact of moore's law: Evidence from when it faltered.

Tobin, J. and Brainard, W. C. (1976). Asset markets and the cost of capital. Technical report, Cowles Foundation for Research in Economics, Yale University.

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

White, B. W. and Rosenblatt, F. (1963). Principles of neurodynamics: Perceptrons and the theory of brain mechanisms. *The American Journal of Psychology*, 76(4):705.

Wildasin, D. E. (1984). The q theory of investment with many capital goods. 74(1):203–210.

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

Yang, S. and Brynjolfsson, E. (2001). Intangible assets and growth accounting: Evidence from computer investments. *Unpublished paper. MIT*, 85(61):28.

Zhang, D., Mishra, S., Brynjolfsson, E., Etchemendy, J., Ganguli, D., Grosz, B., Lyons, T., Manyika, J., Niebles, J. C., Sellitto, M., et al. (2021). The AI index 2021 annual report. *arXiv preprint arXiv:2103.06312*.

Zolas, N., Kroff, Z., Brynjolfsson, E., McElheran, K., Beede, D. N., Buffington, C., Goldschlag, N., Foster, L., and Dinlersoz, E. (2021). Advanced technologies adoption and use by US firms: Evidence from the annual business survey. Technical report, National Bureau of Economic Research.

Zyontz, S. (2018). Running with (CRISPR) scissors: Tool specialists and innovation with new technologies.

10 Appendix

10.1 Data Supplement

Compustat Data: The Compustat/Capital IQ data used come from the Compustat North America database accessed via Wharton Research Data Services (WRDS). For the primary outcome variable of Market Value, I calculate the total assets (*at*) less the common equity book value (*ceq*) and add back the market value of the common equity (*prcc_ c* times *csho*) all at the quarter time period. Other variables used in the analysis include gross value of property, plant, and equipment as an additional asset measure *ppegt*, book value of intangible assets *intan*, goodwill *gdwl*, cash and equivalents *che*, other assets *oa*, and employee count *emp* where populated (otherwise it is inferred in a procedure described below). Industry identifiers are from Compustat's NAICS code field, and the primary key for a firm as an observational unit is the Compustat *gvkey* and ticker symbol.

Occupation Count Measures: Some specifications include a measure of the overall human capital of a firm within a given time period, and specifications for overall engineering value include counts or wage measures by employee function within firms. These measures can be normalized against the BLS-OES data and Compustat/Capital IQ. For firm-level aggregate employment data, I use the Compustat/Capital IQ North America database value of EMP (employee counts). In the case that the EMP value is missing or erroneous, I substitute the predicted value of EMP from a linear model trained on known EMP values of the following form[34]:

$$E\hat{M}P_{it} = \beta_0 + \beta_1 LI_{it} + \beta_3 TA_{it} + \gamma_{jt} + \nu_{it} \tag{22}$$

	(1)	(2)	(3)	(4)
Total Assets	0.00004	0.00004 ∗ ∗	0.00004 ∗ ∗∗	0.00003 ∗ ∗∗
	(0.0000278)	(0.0000187)	(0.0000146)	(0.0000083)
LinkedIn Worker Count	0.00200 ∗ ∗∗	0.00192 ∗ ∗∗	0.00190 ∗ ∗∗	0.00145 ∗ ∗∗
	(0.0005717)	(0.0004699)	(0.0004729)	(0.0003746)
NAICS2-Year FE	✓			
NAICS3-Year FE		✓		
NAICS4-Year FE			✓	
Firm and Year FE				✓
R^2	.3243493	.3897643	.4217879	.9414679
N	53,699	53,607	52,767	53,657
Note:				*p<0.1; **p<0.05; ***p<0.01

Table 10: Prediction of Compustat Employment with LinkedIn Counts

Table Notes: Robust standard errors in parentheses. This reports the results of a regression of Compustat's employee variable (EMP) on the total assets of the firm and the LinkedIn employee count, along with a fixed effect for industries and firms. Since EMP is in thousands of employees, the interpretation of the coefficient on the LinkedIn worker count is how many employees exist per LinkedIn profile. This plot also appears in Tambe et al. (2020a).

The predicted EMP for firm i in year t is a function of the intercept, the LinkedIn total count for that firm in that year (LI_{it}), the total assets of the firm TA_{it}, a fixed effect for that industry-year combination γ_{jt}, and an error term ν_{it}. With knowledge of the total firm-year varying employment, the industry classification (3-Digit NAICS Code), the LinkedIn employment counts by LinkedIn occupational category, and the industry-level employment composition according to the BLS-OES, I build a firm-year-occupation-level coverage ratio for all of the publicly traded firms in Compustat/Capital IQ. Whereas omitting the occupational coverage differences within firm implicitly assumes all workers in the same firm face the same incentives to post information to their profile, this adjustment assumes that all workers with the same occupation in the same firm in the same year are subject to similar data supply incentives. Firm-level differences and year-level differences in coverage are even more substantial and handled by this procedure. Meanwhile this adjustment does make a potentially significant assumption that workers employed by U.S. publicly traded firms but working elsewhere are employed in similar proportions to the BLS-OES industrial occupational employment shares. The appendix has the regression results for equation 22 in Table 10. Typically firms have about 1.9 times as many employees as are available on LinkedIn, controlling for the asset base size and industry-year (using specification 3 from that table).

In detail, first I take the occupational employment shares by industry-year from the BLS-OES. I then calculate the industry-employment shares by industry from Compustat using either EMP or predicted EMP from equation 22. Re-weighting the BLS-OES occupation-industry-year shares by

[34] Prediction accuracy gains from models with higher complexity (e.g. tree-based models or support vector machines) were relatively small.

the Compustat industry-year shares and summing by occupation yields the Compustat occupation-year shares. These Compustat occupation-year shares are multiplied by total Compustat employment (emp or predicted emp) to get the total Compustat employment by occupation-year. The total employment by occupation in publicly traded firms on LinkedIn is compared to this Compustat employment by occupation value to get a job-year-level coverage value μ_{jt}^{job} for the proportion of Compustat employment in job j and year t captured on LinkedIn. The total LinkedIn count in year t at the firm i is then divided by the total Compustat employment in that firm to get μ_{it}^{firm}, the firm-year coverage ratio. Multiplying these two factors is analogous to flipping two biased coins – one for if the worker in firm i is captured by Compustat and LinkedIn, and another for if the worker with job j is on Compustat and LinkedIn. Since these coverages will double-count the employment weighted average coverage ratio by firm $\bar{\mu}_{jt}^{firm}$, I divide that out such that total adjusted LinkedIn employment is equal to total Compustat employment. The relatively simple normalization function to convert observed LinkedIn occupation-firm-year counts into BLS-OES-Compustat standard occupation-firm-year counts is as follows:

$$LI_{ijt} = \frac{(\mu_{it}^{firm}\mu_{jt}^{job})}{\bar{\mu}_{jt}^{firm}} CompustatEmployment_{it} \tag{23}$$

The end result is Compustat-BLS-OES-consistent firm-year-occupation employment coverage ratios. LinkedIn defines Engineering, Information Technology, and Research as separate functional areas within a firm. When members submit their profile information, they are additionally classified into a given functional area. Occupations are distributed across these different domains, not always into the same functional area of the company. Software engineers are most frequently included in the Engineering category (as are most occupations with "engineer" in the title), but may also be categorized in Information Technology. I calculate the total employee counts in each of these different categories. The normalized counts of workers are taken as the output of the adjustment represented in Data Appendix equation 23. Those employee counts are multiplied by their BLS-OES wage in the relevant respective year to construct the wage bill variables for analysis of overall engineering talent. The figures below detail coverage ratios over time by occupation grouping.

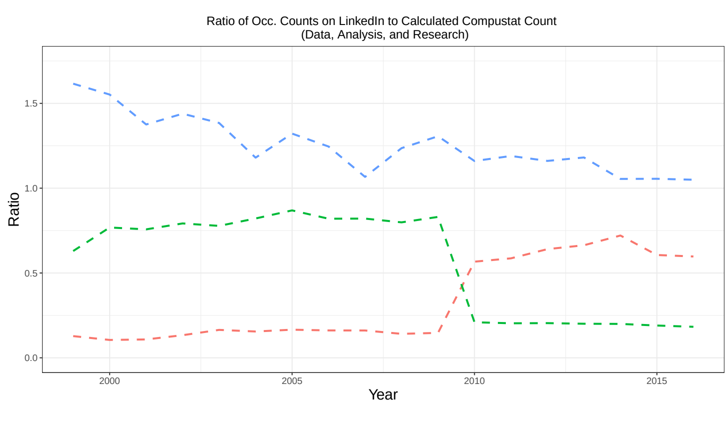

Figure 11: Data-related Occupation LinkedIn Coverage

Figure Notes: This figure plots the LinkedIn count of Research Fellows, Data Analysts, and Database Engineers in Compustat firms over time relative to the expected amounts computed from the procedure in the appendix of Tambe et al. (2020a).

Additional LinkedIn Coverage Plots: Below are plots showing the firm-level histogram over

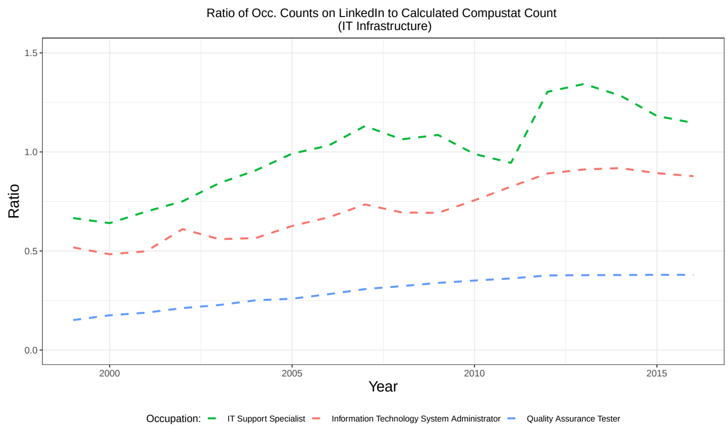

Figure 12: IT-related Occupation LinkedIn Coverage

Figure Notes: This figure plots the LinkedIn count of IT support staff, Systems Administrators, and Quality Assurance employees in Compustat firms over time relative to the expected amounts computed from the procedure in the appendix of Tambe et al. (2020a).

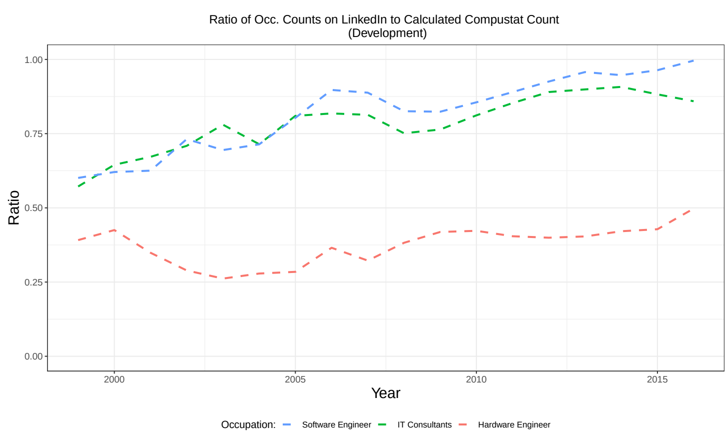

Figure 13: Software Engineering-related Occupation LinkedIn Coverage

Figure Notes: This figure plots the LinkedIn count of software engineers, IT consultants, and hardware engineers in Compustat firms over time relative to the expected amounts computed from the procedure in the appendix of Tambe et al. (2020a).

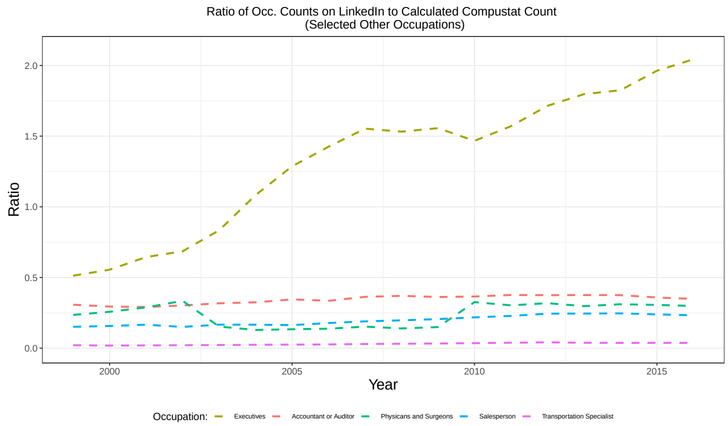

Figure 14: Selected Other Occupation LinkedIn Coverage

Figure Notes: This figure plots the LinkedIn count of executives, accountants, doctors, salespeople, and transportation workers in Compustat firms over time relative to the expected amounts computed from the procedure in the appendix of Tambe et al. (2020a).

coverage rates over time. The comparison is the total LinkedIn employees counted in public firms compared to the EMP variable from Compustat.

10.2 Supplementary Tables for Event Study Coefficients and Other Robustness Checks

Figure 15: LinkedIn Firm Coverage by Year

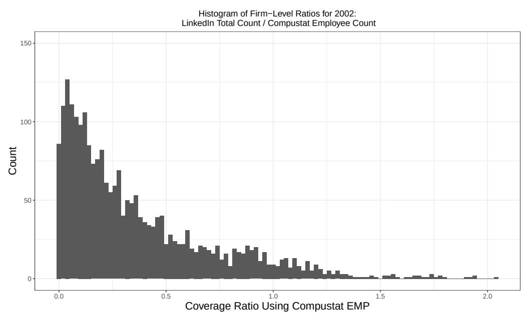

(a) 2002

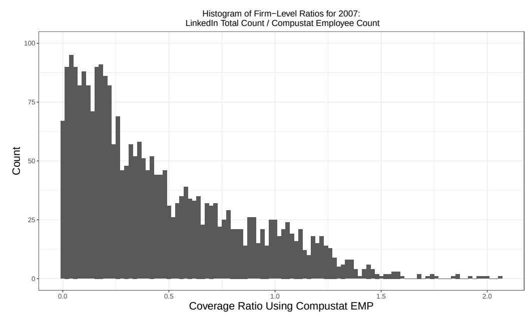

(b) 2007

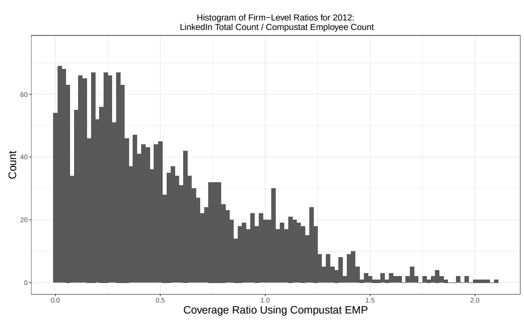

(c) 2012

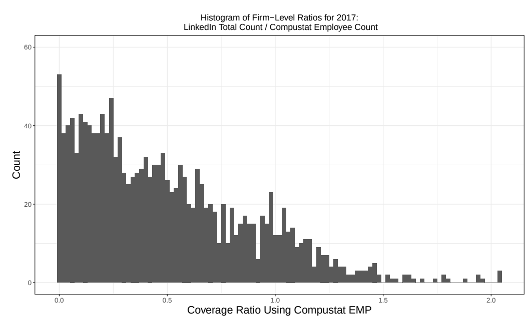

(d) 2017

Figure Notes: These charts show the histogram of firm employment coverage ratios for selected years. The coverage ratio is the total count of resume records on LinkedIn divided by the Compustat EMP database item.

Event Study Coefficients	(1) Full Sample	(2) Dropping Top Quintile
Log(AI Index)	-2,236**	-913.7**
	(960.3)	(427.7)
Log(Data Science Index)	76.15	39.15
	(209.9)	(99.07)
Log(AI Index)-Period 1	0	0
	(0)	(0)
Log(AI Index)-Period 2	284.2*	310.9***
	(171.7)	(115.0)
Log(AI Index)-Period 3	703.0***	320.4**
	(263.6)	(159.5)
Log(AI Index)-Period 4	984.0***	584.8**
	(326.4)	(233.0)
Log(AI Index)-Period 5	975.4**	760.1**
	(404.3)	(305.3)
Log(AI Index)-Period 6	1,106**	729.2**
	(449.1)	(296.0)
Log(AI Index)-Period 7	497.3	491.3
	(465.1)	(335.4)
Log(AI Index)-Period 8	1,107*	770.7**
	(583.1)	(364.0)
Log(AI Index)-Period 9	943.1	576.5
	(577.1)	(376.2)
Log(AI Index)-Period 10	1,085*	607.1
	(594.2)	(399.5)
Log(AI Index)-Period 11	1,540**	704.7*
	(694.7)	(386.7)
Log(AI Index)-Period 12	1,817***	743.9*
	(672.0)	(384.6)
Log(AI Index)-Period 13	2,535***	970.8**
	(836.8)	(407.4)
Log(AI Index)-Period 14	2,752***	1,043**
	(891.1)	(415.4)
Log(AI Index)-Period 15	3,198***	942.9**
	(936.8)	(401.3)
Observations	23,944	19,221
Firm and Quarter FE	✓	✓

Table 11: Event Study Coefficients with and without the Top Quintile of AI Firms

Table Notes: *** p<0.01, ** p<0.05, * p<0.1. Standard errors clustered by firm for firm fixed effects in parentheses. The second column performs the event study in equation **??** but drops the top quintile of AI using firms. Skill index and control variable estimate values other than AI and Data Science are omitted for display purposes, but very similar to values in other specifications. Period 1 is the reference period and has coefficients equal to 0. Specification and included controls are based on column (5) of Table 5.

Event Study by AI Quintile	(1) Log Market Value	(2) Log Market Value	(3) Market Value	(4) Market Value
Quintile 3x2	-0.0375	0.00337	-395.4	9.368
	(0.0344)	(0.0127)	(414.7)	(148.0)
Quintile 3x3	-0.00299	0.00617	-764.3	-50.32
	(0.0282)	(0.0192)	(526.2)	(122.0)
Quintile 3x4	-0.0340	0.00893	-463.7	132.8
	(0.0275)	(0.0241)	(487.2)	(276.6)
Quintile 3x5	-0.0368	0.00422	-95.23	55.29
	(0.0352)	(0.0267)	(541.9)	(263.8)
Quintile 3x6	-0.0438	-0.00149	-622.3	-105.7
	(0.0381)	(0.0269)	(558.0)	(284.6)
Quintile 3x7	-0.0256	-0.0259	-1,182**	-545.3
	(0.0422)	(0.0338)	(561.1)	(575.8)
Quintile 3x8	-0.00394	0.0248	-993.6	-204.2
	(0.0526)	(0.0358)	(607.9)	(452.4)
Quintile 3x9	0.00833	0.0248	-383.7	-89.17
	(0.0495)	(0.0371)	(508.9)	(333.7)
Quintile 3x10	0.0167	0.0440	-715.2*	-115.0
	(0.0561)	(0.0391)	(406.4)	(319.9)
Quintile 3x11	0.0140	0.0430	-396.2	-149.3
	(0.0602)	(0.0404)	(372.4)	(320.7)
Quintile 3x12	0.00283	0.0448	-1,188*	-73.99
	(0.0644)	(0.0421)	(619.6)	(274.1)
Quintile 3x13	0.00329	0.0408	-205.9	48.33
	(0.0617)	(0.0450)	(407.6)	(337.3)
Quintile 3x14	0.00770	0.0326	-92.59	-66.15
	(0.0671)	(0.0486)	(526.6)	(305.5)
Quintile 4x2	-0.00154	0.000159	332.0	339.3**
	(0.0139)	(0.00901)	(465.1)	(133.0)
Quintile 4x3	-0.00207	-0.000620	414.3	365.9**
	(0.0203)	(0.0132)	(464.5)	(171.2)
Quintile 4x4	-0.00290	0.0151	1,057*	876.9***
	(0.0238)	(0.0156)	(628.8)	(253.5)
Quintile 4x5	-0.00111	0.0127	601.1	744.5***
	(0.0249)	(0.0171)	(589.5)	(250.9)
Quintile 4x6	-0.00954	0.0162	507.1	906.1***
	(0.0247)	(0.0182)	(827.4)	(282.7)
Quintile 4x7	-0.00246	0.00922	-975.3**	129.2
	(0.0249)	(0.0213)	(408.6)	(248.5)
Quintile 4x8	-0.0277	0.00936	-624.5	221.4
	(0.0336)	(0.0219)	(542.3)	(320.6)
Quintile 4x9	-0.0415	-0.0135	82.13	93.23
	(0.0302)	(0.0236)	(677.7)	(412.1)
Quintile 4x10	-0.0345	-0.00976	442.1	235.7
	(0.0319)	(0.0229)	(400.7)	(362.6)
Quintile 4x11	-0.0263	-0.00626	651.8	401.6
	(0.0308)	(0.0233)	(587.6)	(362.2)
Quintile 4x12	-0.0496	0.00136	-624.3	643.4
	(0.0364)	(0.0239)	(503.3)	(394.6)
Quintile 4x13	-0.0671*	-0.00809	76.81	779.8*
	(0.0353)	(0.0254)	(348.5)	(467.3)
Quintile 4x14	-0.0810**	-0.0262	879.4**	771.9*
	(0.0401)	(0.0270)	(380.9)	(451.7)
Quintile 5x2	0.00208	0.00453	1,126	761.3
	(0.0134)	(0.00651)	(1,177)	(469.1)
Quintile 5x3	0.0446*	0.0260***	1,588	1,580**
	(0.0241)	(0.00917)	(1,845)	(736.0)
Quintile 5x4	0.0238	0.0281**	2,737*	2,723***
	(0.0250)	(0.0116)	(1,472)	(836.1)
Quintile 5x5	0.0291	0.0314**	406.2	2,010***
	(0.0268)	(0.0135)	(1,412)	(729.0)
Quintile 5x6	0.00653	0.0276*	171.9	2,096**
	(0.0260)	(0.0146)	(1,564)	(849.2)
Quintile 5x7	0.0358	0.0361**	-2,732**	-551.4
	(0.0279)	(0.0174)	(1,083)	(808.3)
Quintile 5x8	0.0261	0.0333*	-3,701***	-62.59
	(0.0214)	(0.0190)	(1,405)	(826.4)
Quintile 5x9	0.0204	0.0176	-824.0	847.2
	(0.0236)	(0.0198)	(1,303)	(804.5)
Quintile 5x10	0.0288	0.0197	3,635***	2,927***
	(0.0229)	(0.0199)	(1,151)	(1,114)
Quintile 5x11	0.0121	0.00923	3,703**	3,847***
	(0.0279)	(0.0210)	(1,425)	(1,020)
Quintile 5x12	-0.0154	0.0218	-2,392	3,946***
	(0.0272)	(0.0221)	(1,507)	(1,063)
Quintile 5x13	-0.0253	0.0305	1,607*	5,689***
	(0.0303)	(0.0233)	(949.2)	(1,235)
Quintile 5x14	-0.0352	0.0187	4,549**	7,721***
	(0.0291)	(0.0243)	(1,954)	(1,413)
Constant	0.443***	2.880***	-2,949***	-10,614*
	(0.0915)	(0.387)	(989.0)	(5,434)
Observations	21,129	21,841	22,736	23,499
Firm and Quarter FE	X	✓	X	✓
Industry-Quarter FE	✓	X	✓	X

Table 12: Event Study Coefficients by Quintile of AI Index

Table Notes: *** p<0.01, ** p<0.05, * p<0.1. Standard errors clustered by firm for firm fixed effects and by 4-Digit NAICS for industry-time fixed effects specifications in parentheses. Coefficients reflect event study specifications, with the first period as the reference period and non-AI-using firms as the reference group, and all skill indices in Table 5 included. The coefficients other than quintiles interacted with time dummies are omitted for display purposes, but very similar to values in other specifications. Specification and included controls are based on column (5) of Table 5.

AI Skills Difference-in-Difference	(1) Market Value	(2) Market Value	(3) Market Value	(4) Market Value	(5) Market Value	(6) Market Value	(7) Market Value
Total Assets	1.072***	1.072***	1.072***	1.072***	1.072***	0.870***	0.870***
	(0.0440)	(0.0440)	(0.0440)	(0.0440)	(0.0440)	(0.0498)	(0.0502)
Log(Human Capital)	1,380***	1,371***	1,386***	1,390***	1,404***	1,020***	1,036***
	(505.8)	(501.1)	(504.6)	(505.3)	(504.7)	(333.3)	(336.0)
Log(AI Index)	-990.5*	-1,011*	-968.0*	-965.9*	-959.9*	-784.0**	-796.5**
	(552.3)	(559.1)	(549.0)	(548.3)	(545.6)	(380.2)	(381.5)
Log(AI Index)xPostTF	1,086***	1,081***	1,090***	1,090***	1,095***	781.0***	791.2***
	(420.0)	(418.7)	(420.7)	(420.9)	(422.5)	(268.7)	(269.9)
Log(Data Science Index)		188.9	263.9	274.0	203.3	128.5	112.7
		(211.4)	(217.4)	(217.7)	(198.2)	(163.5)	(165.6)
Log(Cloud Computing Index)			-337.6*	-330.5*	-361.5*	-290.0**	-288.8**
			(179.5)	(178.7)	(187.0)	(137.9)	(138.4)
Log(Data Storage)				-8,258	-7,930	-3,229	-3,101
				(7,674)	(7,711)	(5,832)	(5,833)
Log(Digital Literacy Index)					-197.9	-121.1	-109.5
					(168.8)	(149.0)	(149.4)
Log(Management Index)					546.4	446.4	444.6
					(412.4)	(343.0)	(341.9)
Log(Advertising)					-11.96	-3.741	-3.026
					(132.0)	(107.2)	(107.3)
Lagged MV						0.253***	0.253***
						(0.0631)	(0.0639)
Lagged MV Growth							-698.0*
							(363.5)
2 Qtr Lagged MV Growth							-260.1
							(251.9)
Constant	-8,019	-8,639	-8,351	-8,281	-9,979	-9,159*	-9,324*
	(5,246)	(5,565)	(5,501)	(5,491)	(6,947)	(5,010)	(5,002)
Observations	23,944	23,944	23,944	23,944	23,944	23,904	23,820
Firm and Quarter FE	✓	✓	✓	✓	✓	✓	✓

Table 13: TensorFlow Difference-in-Difference on Unbalanced Panel

Table Notes: *** $p<0.01$, ** $p<0.05$, * $p<0.1$. Standard errors clustered by firm in parentheses. The table reports results from estimating versions of equation 18 on an unbalanced panel of firms from the beginning of 2014 to the end of 2017. The logged indices, including human capital, are actually $log(x+1)$ in all cases. Column (1) reports a specification with no additional adjustments other than the fixed effects, total assets, and human capital. Columns (2) to (5) add in different skills indices, including advertising (in case AI is mostly useful for ad tech). Columns (6) and (7) include controls for lagged market value (1 quarter) and lagged market value growth (1 and 2 quarters) respectively in order to further control for possible trends in market value like AI hype or momentum in investment. Overall the AI index x Post-TF estimates are relatively stable between $7.8 and $10.9 million more market value per 1 percent increase in AI skills in the post period. In other words, TensorFlow appears to cause proportionate increase of $780 million to $1 billion per 100 percent increase of AI talent in companies. The unbalanced panel is broadly consistent with the balanced one.

AI Skills Difference-in-Difference	(1) Market Value	(2) Market Value	(3) Market Value	(4) Market Value	(5) Market Value	(6) Market Value	(7) Market Value
Total Assets	1.089***	1.089***	1.090***	1.090***	1.090***	0.878***	0.876***
	(0.0545)	(0.0545)	(0.0545)	(0.0545)	(0.0544)	(0.0533)	(0.0538)
Log(Human Capital)	1,602***	1,594***	1,613***	1,617***	1,640***	1,151***	1,153***
	(587.8)	(583.1)	(587.0)	(587.5)	(588.6)	(381.1)	(381.5)
Log(AI Index)	-1,049*	-1,070*	-1,017*	-1,014*	-1,007*	-816.7**	-822.2**
	(593.0)	(600.0)	(588.9)	(588.2)	(584.9)	(398.4)	(398.6)
Log(AI Index)xPostTF	1,106**	1,101**	1,111**	1,112**	1,118**	751.1***	755.5***
	(434.7)	(433.5)	(435.5)	(435.6)	(437.8)	(268.2)	(268.2)
Log(Data Science Index)		186.3	278.1	287.0	183.3	122.7	111.7
		(229.5)	(236.1)	(236.0)	(215.5)	(178.2)	(179.6)
Log(Cloud Computing Index)			-403.1**	-394.4**	-436.7**	-345.4**	-345.8**
			(190.4)	(189.9)	(200.1)	(143.1)	(143.5)
Log(Data Storage Index)				-9,203	-9,093	-4,123	-4,135
				(7,700)	(7,773)	(5,766)	(5,787)
Log(Digital Literacy Index)					-228.3	-127.9	-123.2
					(179.4)	(166.0)	(165.9)
Log(Management Index)					681.4	548.2	534.1
					(470.4)	(398.7)	(396.8)
Log(Advertising Index)					4.805	-5.366	0.575
					(140.9)	(114.5)	(114.3)
Lagged MV						0.270***	0.272***
						(0.0662)	(0.0670)
Lagged MV Growth							-809.5*
							(419.1)
2 Qtr Lagged MV Growth							-295.3
							(291.3)
Constant	-10,562*	-11,191*	-10,863*	-10,779*	-13,039	-11,608**	-11,537**
	(6,173)	(6,533)	(6,464)	(6,456)	(8,122)	(5,614)	(5,570)
Observations	21,275	21,275	21,275	21,275	21,275	21,273	21,260
Firm and Quarter FE	✓	✓	✓	✓	✓	✓	✓

Table 14: TensorFlow Difference-in-Difference for Tobin's q

Table Notes: *** p<0.01, ** p<0.05, * p<0.1. Standard errors clustered by firm in parentheses. The table reports results from estimating versions of equation 18 on a balanced panel of firms from the beginning of 2014 to the end of 2017. The logged indices, including human capital, are actually $log(x+1)$ in all cases. Column (1) reports a specification with no additional adjustments other than the fixed effects, total assets, and human capital. Columns (2) to (5) add in different skills indices, including advertising (in case AI is mostly useful for ad tech). Columns (6) and (7) include controls for lagged market value (1 quarter) and lagged market value growth (1 and 2 quarters) respectively in order to further control for possible trends in market value like AI hype or momentum in investment. Results in this table are substantively similar to the main specifications, but constrain the total asset coefficient to be one.